From Automation to Augmentation

Designing AI Coaches That Amplify Expertise

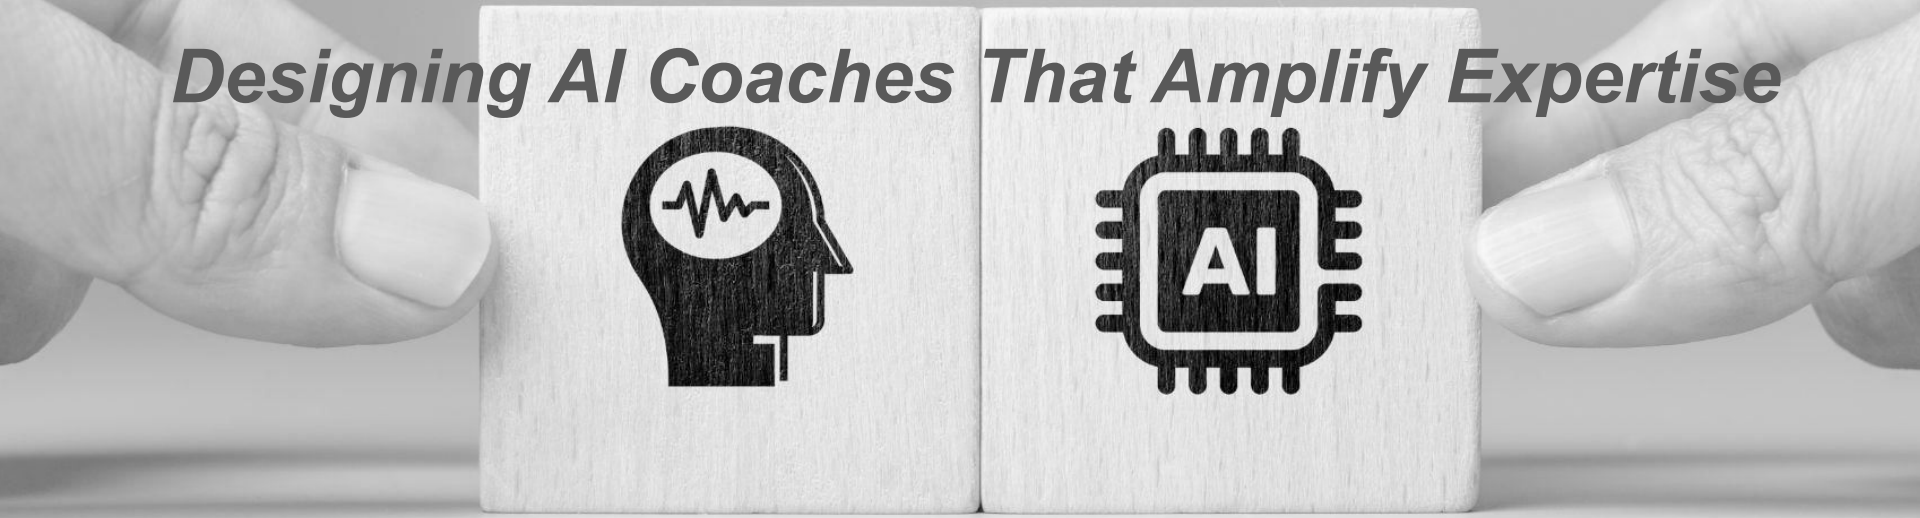

Mike Amundsen
@mamund


Mike Amundsen
@mamund

Approved by Management #26-02

Because Someone Signed Off on This

7 HRS AGO · MIKE AMUNDSEN

Most Popular

VIEW ALL

The Hypermedia Commons We Missed

SEP 23, 2025 · MIKE AMUNDSEN

GenAI's Bezos Mandate? Shopify Just Issued It

APR 8, 2025 · MIKE AMUNDSEN

Don't design workflows, discover them

JAN 19 · MIKE AMUNDSEN

Let’s get started…

Customer Service

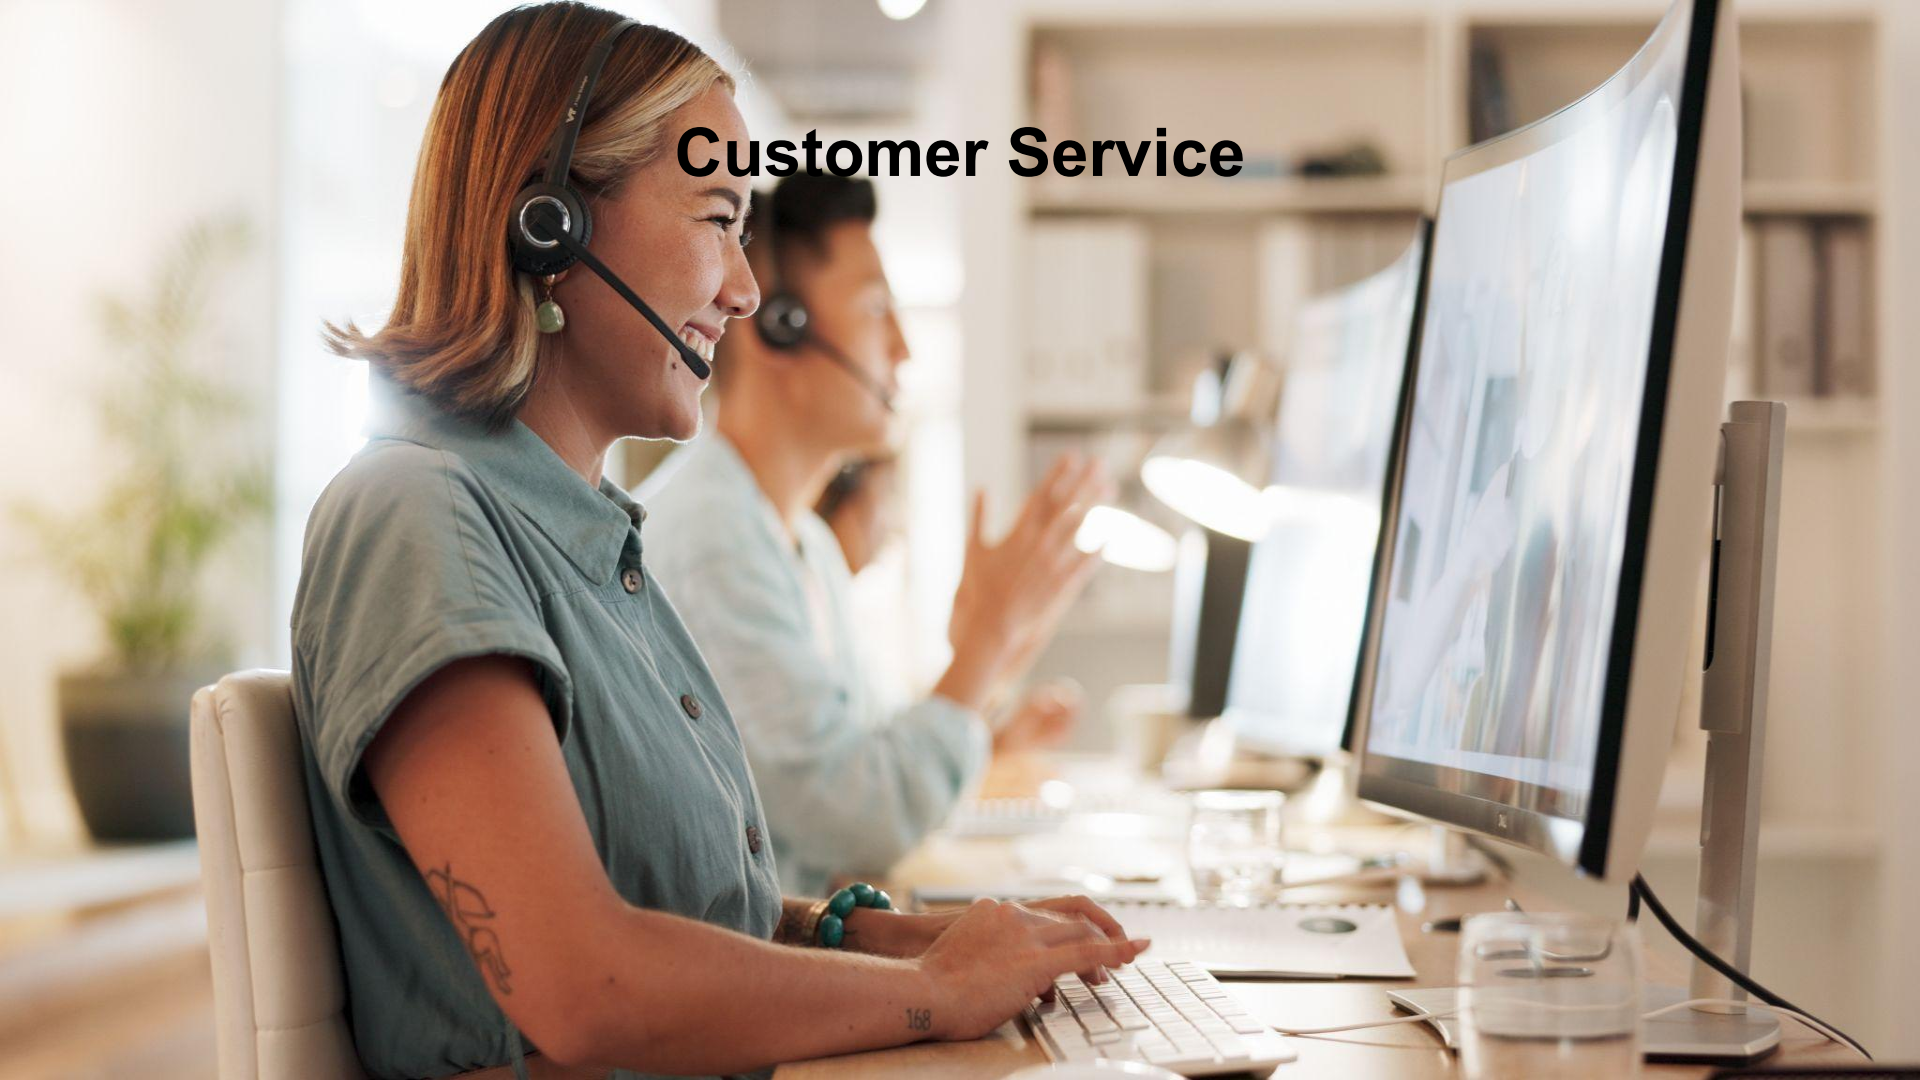

CSR + AI = Trouble


FAST COMPANY

01-12-2026 | MAVEN AGI

Klarna tried to replace its workforce with AI

Here is what every company should learn from it

CSR + AI = Trouble

FAST COMPANY

01-12-2026 | MAVEN AGI

Klarna tried to replace its workforce with AI

Here is what every com

Technology

Salesforce AI faces backlash from customers

Customers angry over the company's switch to AI-driven tech support are saying things like: "It is so infuriating that Salesforce wants to abandon things that currently work."

Oct 4, 2025 10:31 AM EDT

By Vuk Zdinjak

CSR + AI = Trouble

FAST COMPANY

01-12-2026 | MAVEN AGI

Klarna tried to replace its workforce with AI

Here is what every com

Technology

Salesforce AI faces backlash from customers

Customers angry over the company's switch to AI-driven tech support are saying things like: "It is so infuriating that Salesforce wants to abandon things that currently work."

Oct 4, 2025 10:31 AM EDT

By Vuk Zdinjak

TOI The Times of India + Follow 2.4M Followers

How Cursor AI chatbot's Big Mess is a lesson for companies automating their customer service

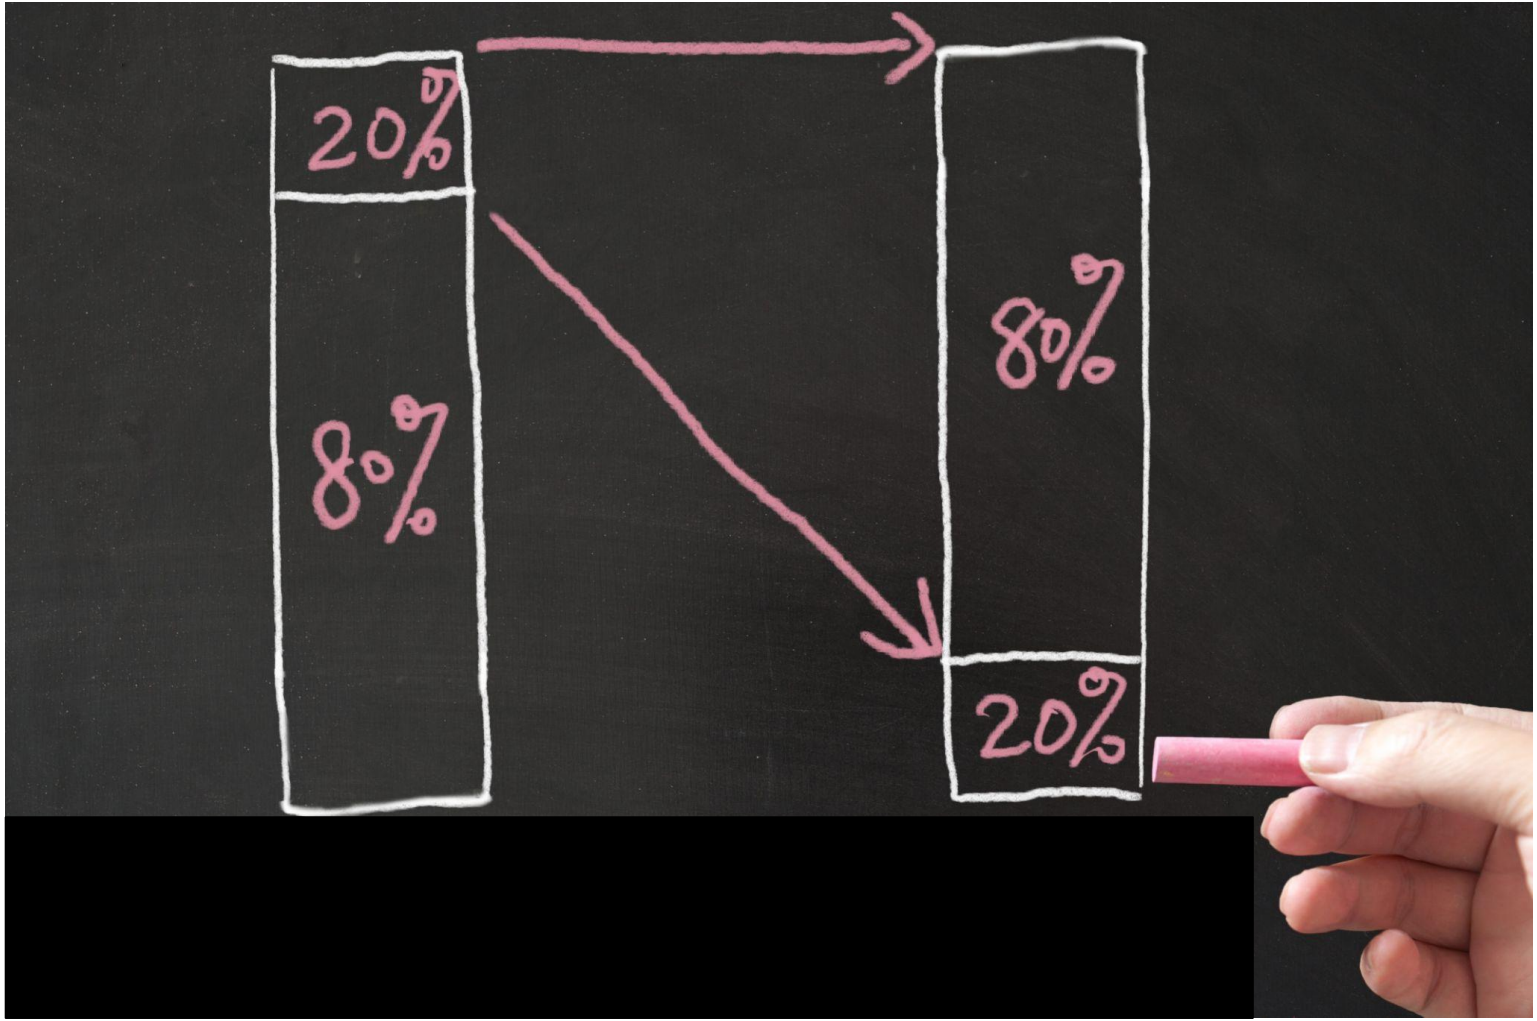
20%
80%
80%
20%

Unintended Consequences

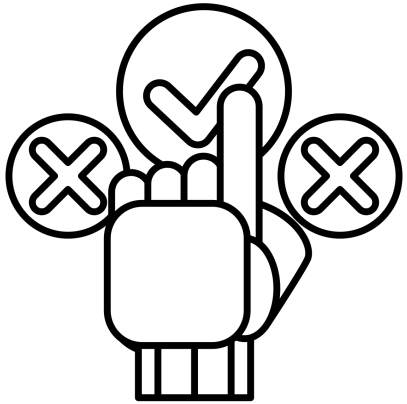

Generative AI and Users

How people use generative AI

- ***As a Generator*** : to get answers
- ***As a Surrogate*** : to get advice
- ***As an Instrument*** : to support/challenge thinking

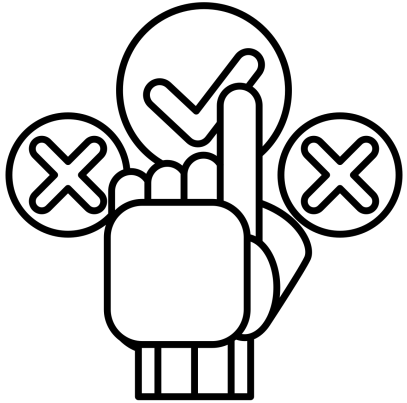

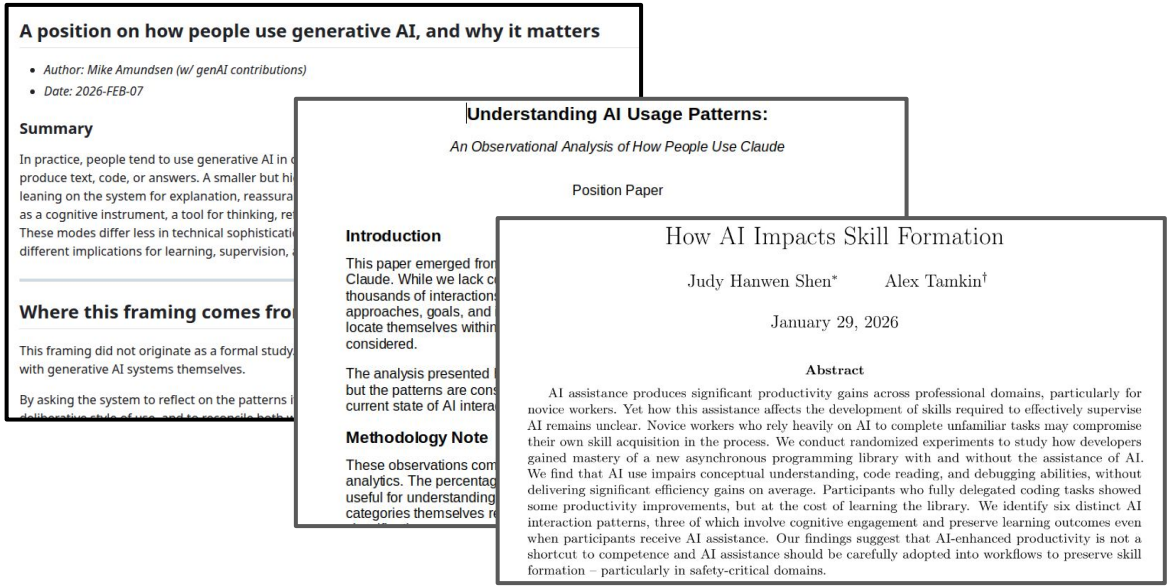

A position on how people use generative AI, and why it matters

- *Author: Mike Amundsen (w/ genAI contributions)*
- *Date: 2026-FEB-07*

Summary

In practice, people tend to use generative AI in
produce text, code, or answers. A smaller but hi
leaning on the system for explanation, reassura
as a cognitive instrument, a tool for thinking, re
These modes differ less in technical sophisticati
different implications for learning, supervision,

Where this framing comes fro

This framing did not originate as a formal study
with generative AI systems themselves.

By asking the system to reflect on the patterns i

Understanding AI Usage Patterns:

An Observational Analysis of How People Use Claude

Position Paper

Introduction

This paper emerged fro
Claude. While we lack c
thousands of interaction
approaches, goals, and
locate themselves within
considered.

The analysis presented
but the patterns are cons
current state of AI intera

Methodology Note

These observations com
analytics. The percentag
useful for understanding
categories themselves r

How AI Impacts Skill Formation

Judy Hanwen Shen* Alex Tamkin†

January 29, 2026

Abstract

AI assistance produces significant productivity gains across professional domains, particularly for novice workers. Yet how this assistance affects the development of skills required to effectively supervise AI remains unclear. Novice workers who rely heavily on AI to complete unfamiliar tasks may compromise their own skill acquisition in the process. We conduct randomized experiments to study how developers gained mastery of a new asynchronous programming library with and without the assistance of AI. We find that AI use impairs conceptual understanding, code reading, and debugging abilities, without delivering significant efficiency gains on average. Participants who fully delegated coding tasks showed some productivity improvements, but at the cost of learning the library. We identify six distinct AI interaction patterns, three of which involve cognitive engagement and preserve learning outcomes even when participants receive AI assistance. Our findings suggest that AI-enhanced productivity is not a shortcut to competence and AI assistance should be carefully adopted into workflows to preserve skill formation – particularly in safety-critical domains.

How people use generative AI

- ***As a Generator*** : to get answers
- ***As a Surrogate*** : to get advice
- ***As an Instrument*** : to support/challenge thinking

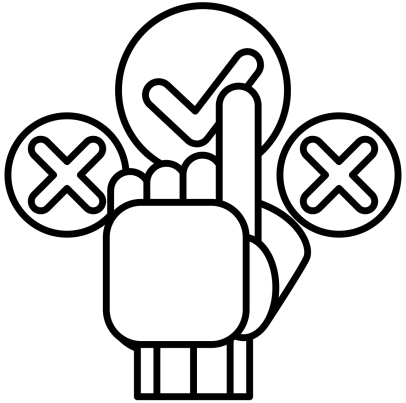

Using AI as an ***Instrument***
is rare

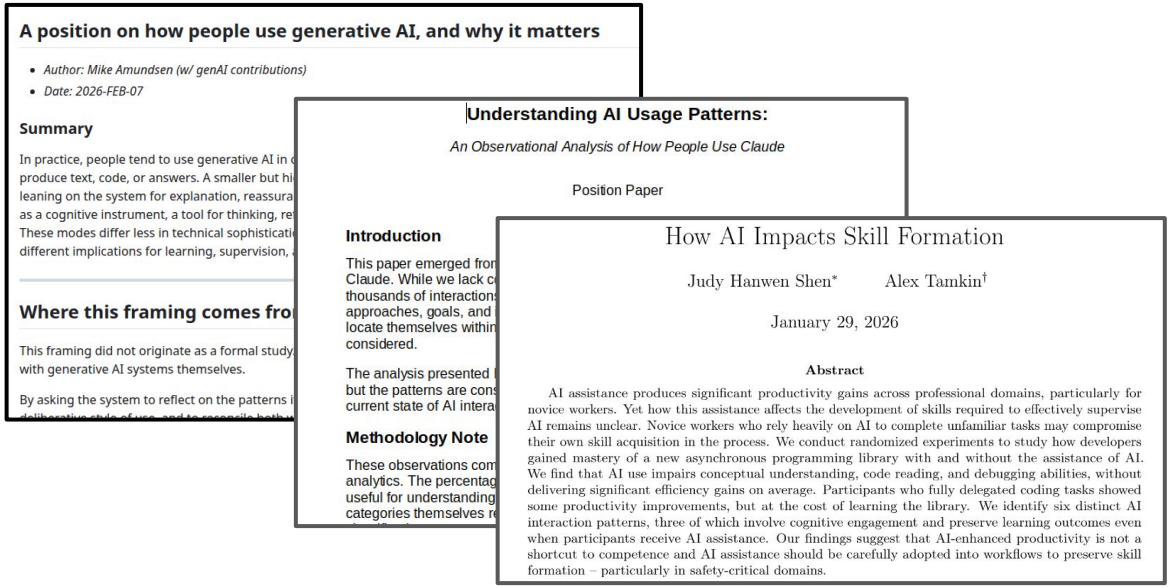

A position on how people use generative AI, and why it matters

- *Author: Mike Amundsen (w/ genAI contributions)*
- *Date: 2026-FEB-07*

Summary

In practice, people tend to use generative AI in
produce text, code, or answers. A smaller but hi
leaning on the system for explanation, reassura
as a cognitive instrument, a tool for thinking, re
These modes differ less in technical sophisticati
different implications for learning, supervision,

Where this framing comes fro

This framing did not originate as a formal study
with generative AI systems themselves.

By asking the system to reflect on the patterns i

Understanding AI Usage Patterns:

An Observational Analysis of How People Use Claude

Position Paper

Introduction

This paper emerged fro
Claude. While we lack c
thousands of interaction
approaches, goals, and
locate themselves within
considered.

The analysis presented
but the patterns are cons
current state of AI intera

Methodology Note

These observations com
analytics. The percentag
useful for understanding
categories themselves r

How AI Impacts Skill Formation

Judy Hanwen Shen* Alex Tamkin†

January 29, 2026

Abstract

AI assistance produces significant productivity gains across professional domains, particularly for novice workers. Yet how this assistance affects the development of skills required to effectively supervise AI remains unclear. Novice workers who rely heavily on AI to complete unfamiliar tasks may compromise their own skill acquisition in the process. We conduct randomized experiments to study how developers gained mastery of a new asynchronous programming library with and without the assistance of AI. We find that AI use impairs conceptual understanding, code reading, and debugging abilities, without delivering significant efficiency gains on average. Participants who fully delegated coding tasks showed some productivity improvements, but at the cost of learning the library. We identify six distinct AI interaction patterns, three of which involve cognitive engagement and preserve learning outcomes even when participants receive AI assistance. Our findings suggest that AI-enhanced productivity is not a shortcut to competence and AI assistance should be carefully adopted into workflows to preserve skill formation – particularly in safety-critical domains.

How people use generative AI

- ***Generator*** : ~60–70%
- ***Surrogate*** : ~20–30%
- ***Instrument*** : ~5–10%

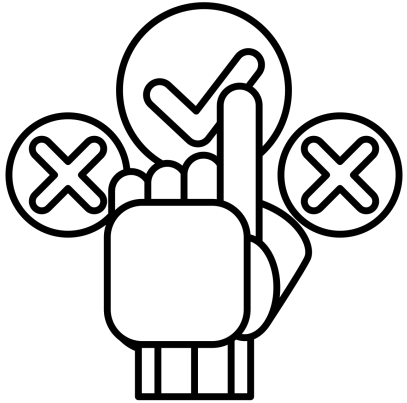

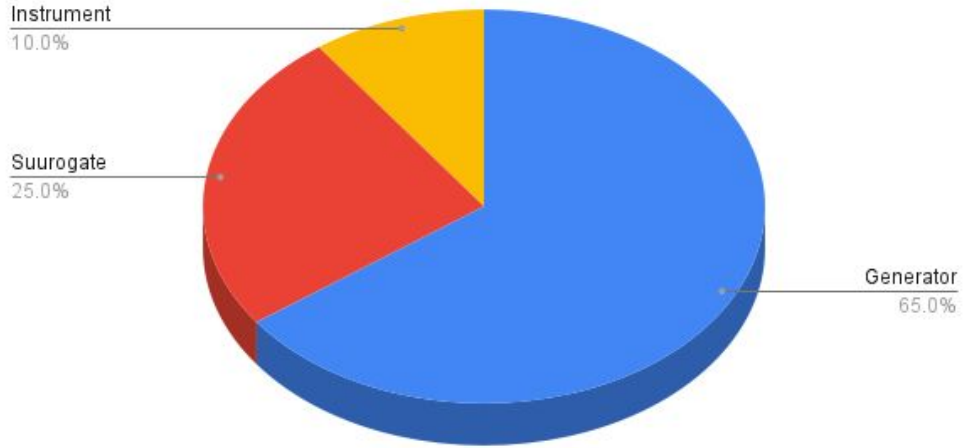

How people use generative AI

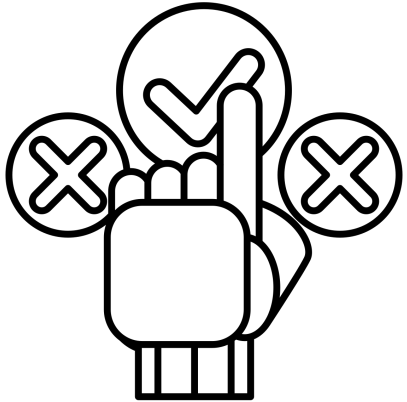

- ***Generator*** : ~60–70%
- ***Surrogate*** : ~20–30%
- ***Instrument*** : ~5–10%

- Judgment shifts
  From humans to machines
- Skill formation weakens
  Disrupting learning culture

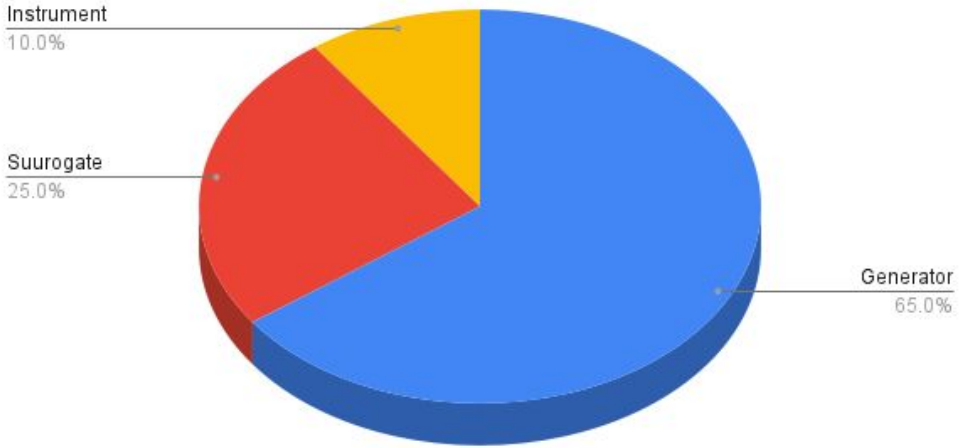

AI Assistance Reduces Persistence and Hurts Independent Performance

Grace Liu
Carnegie Mellon University
gliu2@andrew.cmu.edu

Brian Christian
University of Oxford
brian.christian@psy.ox.ac.uk

Tsvetomira Dumbalska
University of Oxford
tsvetomira.dumbalska@psy.ox.ac.uk

Michiel A. Bakker
Massachusetts Institute of Technology
bakker@mit.edu

Rachit Dubey
University of California, Los Angeles
rdubey@ucla.edu

Abstract

People often optimize for *long-term* goals in collaboration: A mentor or companion doesn't just answer questions, but also scaffolds learning, tracks progress, and prioritizes the other person's growth over immediate results. In contrast, current AI systems are fundamentally *short-sighted* collaborators – optimized for providing instant and complete responses, without ever saying no (unless for safety reasons). What are the consequences of this dynamic? Here, through a series of randomized controlled trials on human-AI interactions ($N = 1,222$), we provide *causal* evidence for two key consequences of AI assistance: reduced persistence and impairment of unassisted performance. Across a variety of tasks, including mathematical reasoning and reading comprehension, we find

What are we optimizing for?

- Speed and efficiency
  vs
- Capability and understanding

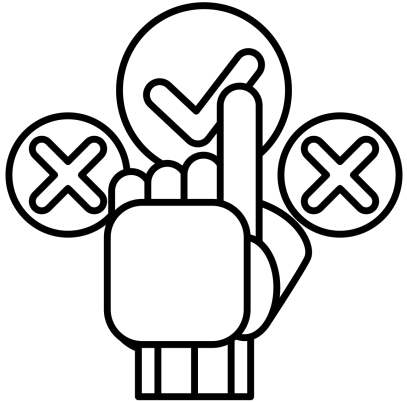

What are we optimizing for?

- Speed and efficiency (automation)
  vs
- Capability and understanding (augmentation)

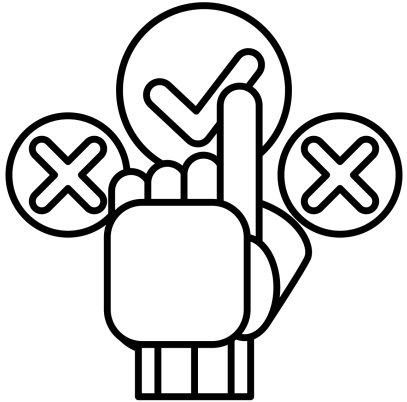

Let's focus on augmentation

What are we optimizing for?

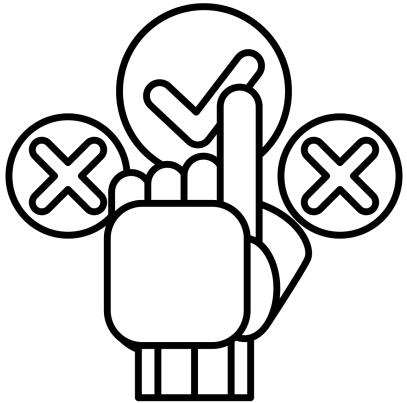

- Speed and efficiency
  vs
- Capability and understanding

Augmentation Legacy:

- V. Bush :
  - Links and trails
- JCR Licklider :
  - Human/computer symbiosis
- Engelbart :
  - Get better at getting better

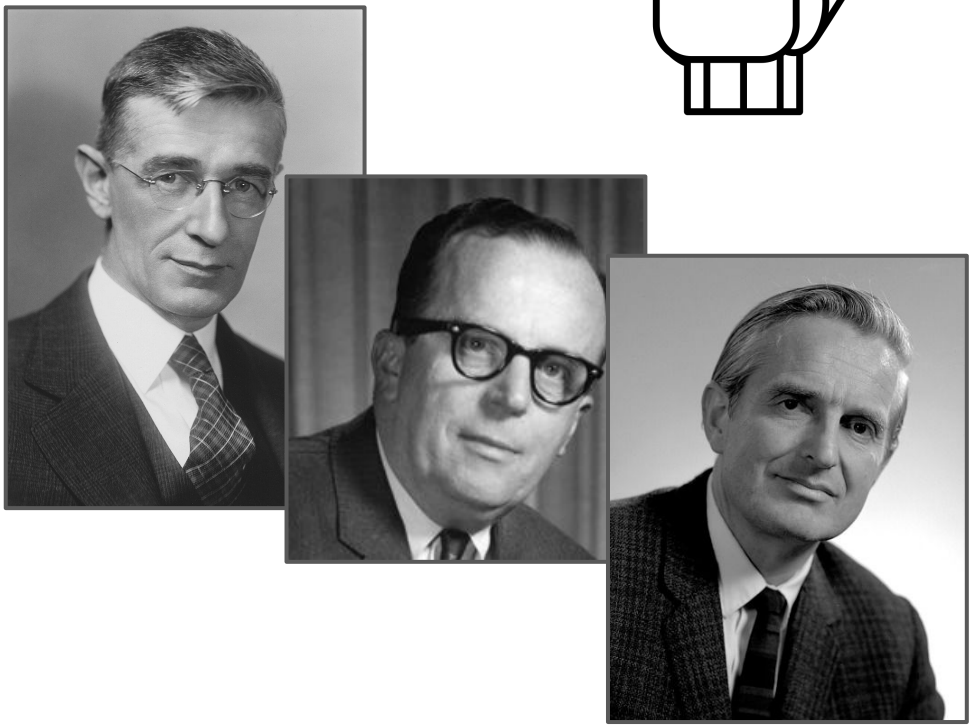

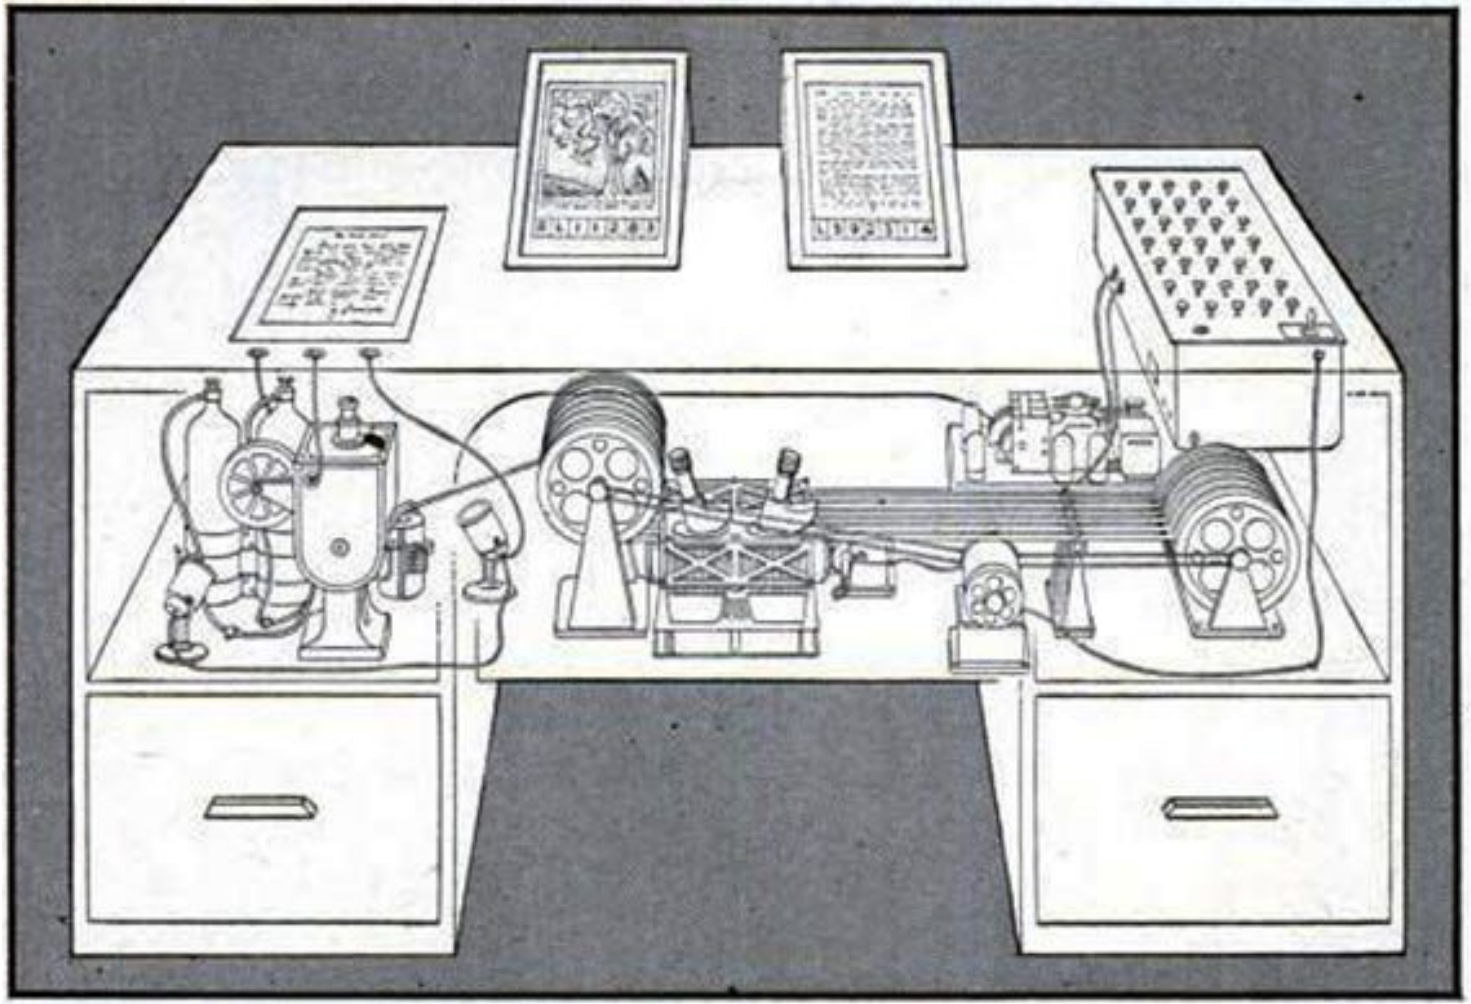

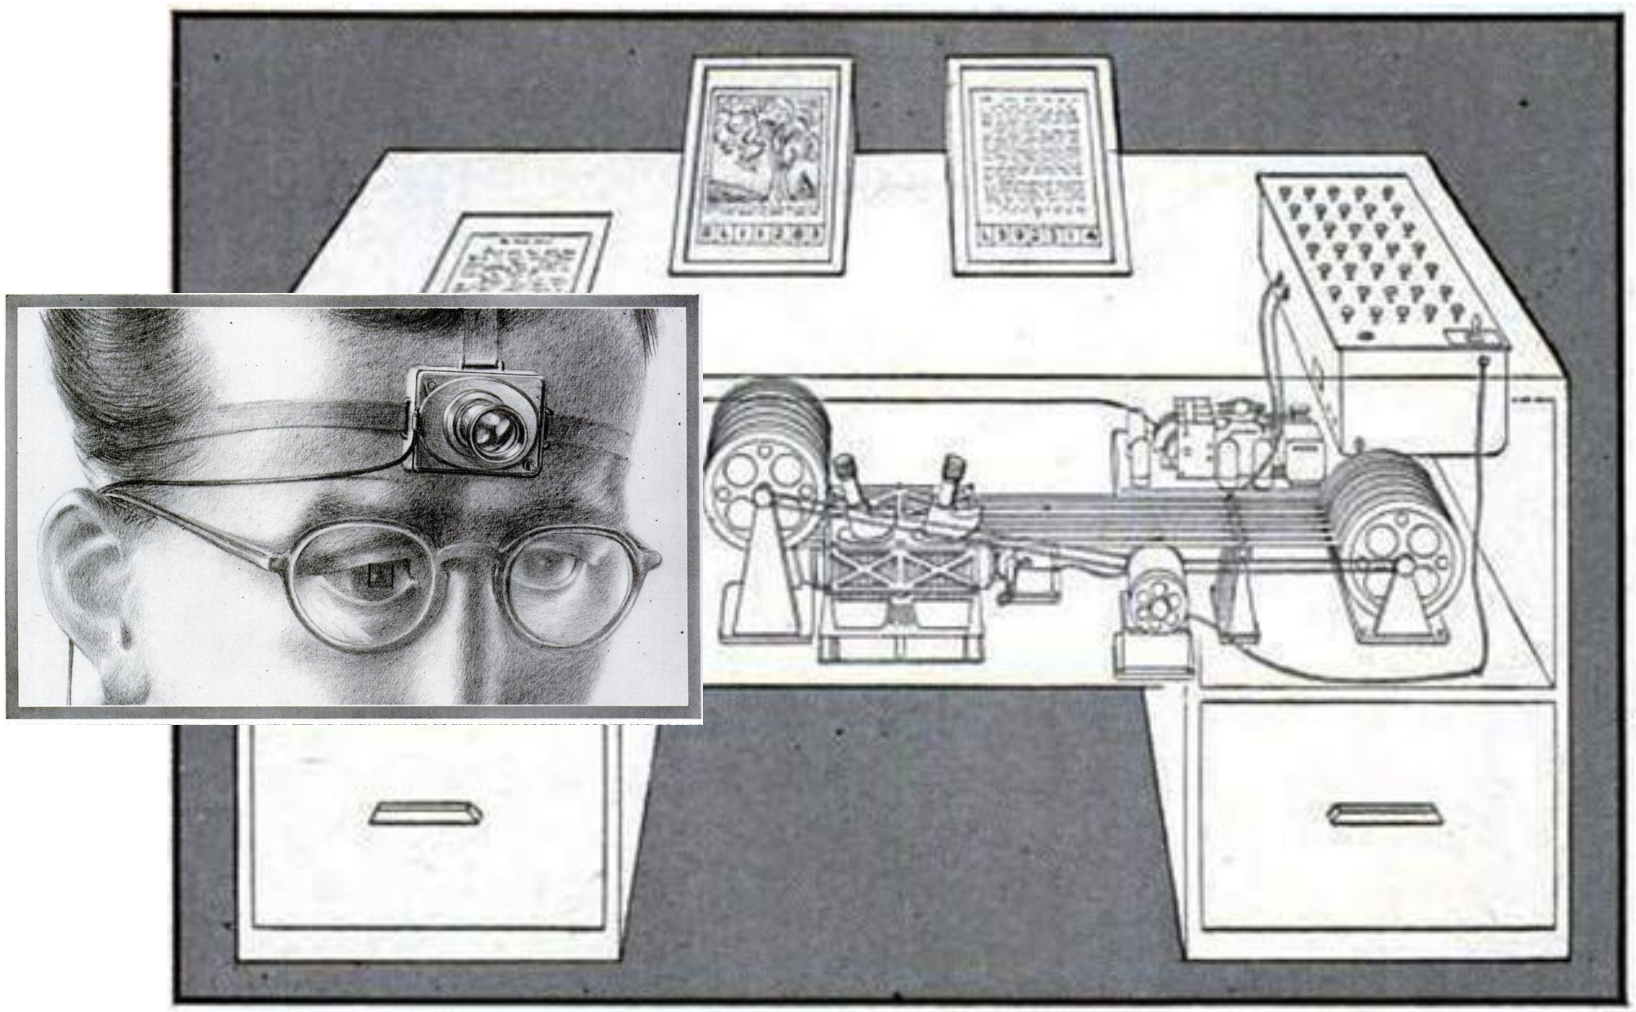

Man-Computer Symbiosis*

J. C. R. LICKLIDER†

Summary—Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers. It will involve very close coupling between the human and the electronic members of the partnership. The main aims are 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In the anticipated symbiotic partnership, men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions in technical and scientific thinking. Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much more effectively than man alone can perform them. Prerequisites for the achievement of the effective, cooperative association include developments in computer time sharing, in memory components, in memory organization, in programming languages, and in input and output equipment.

I. Introduction

will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.

B. Between "Mechanically Extended Man" and "Artificial Intelligence"

As a concept, man-computer symbiosis is different in an important way from what North[2] has called "mechanically extended man." In the man-machine systems of the past, the human operator supplied the initiative, the direction, the integration, and the criterion. The mechanical parts of the systems were mere extensions, first of the human arm, then of the human eye. These systems certainly did not consist of "dissimilar organisms living together . . ." There was only one kind of organism—man—and the rest was there only to help him.

In one sense of course, any man-made system is in-

Man-Computer Symbiosis*

J. C. R. LICKLIDER†

is is an expected develop-tween men and electronic lose coupling between the f the partnership. The main ate formulative thinking as ormulated problems, and 2) perate in making decisions ithout inflexible dependence anticipated symbiotic part-ulate the hypotheses, deter-e evaluations. Computing ork that must be done to decisions in technical and yses indicate that the sym-ellectual operations much more effectively than man alone can perform them. Prerequisites for the achievement of the effective, cooperative association include developments in computer time sharing, in memory components, in memory organization, in programming languages, and in input and output equipment.

I. Introduction

will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.

B. Between "Mechanically Extended Man" and "Artificial Intelligence"

As a concept, man-computer symbiosis is different in an important way from what North[2] has called "mechanically extended man." In the man-machine systems of the past, the human operator supplied the initiative, the direction, the integration, and the criterion. The mechanical parts of the systems were mere extensions, first of the human arm, then of the human eye. These systems certainly did not consist of "dissimilar organisms living together . . ." There was only one kind of organism—man—and the rest was there only to help him.

In one sense of course, any man-made system is in-

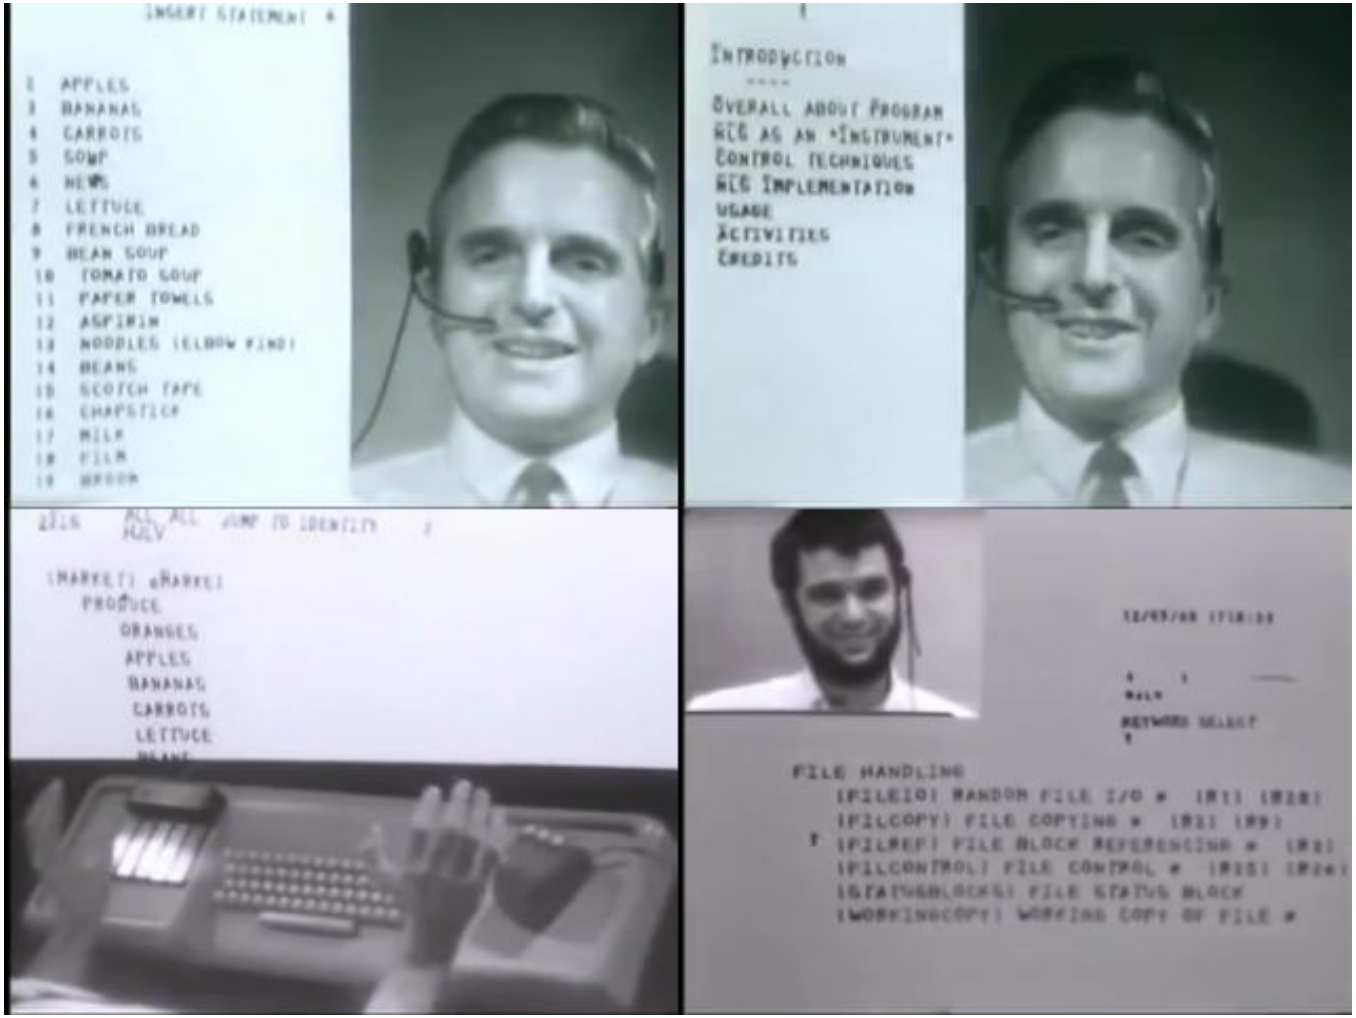
INSERT STATEMENT #
1 APPLES
3 BANANAS
4 CARROTS
5 SOUP
6 MEAT
7 LETTUCE
8 FRENCH BREAD
9 BEAN SOUP
10 TOMATO SOUP
11 PAPER TOWELS
12 ASPIRIN
13 NOODLES (ELBOW KIND)
14 BEANS
15 SCOTCH TAPE
16 CHAPSTICK
17 MILK
18 FILM
19 BROOM
Introduction
OVERALL ABOUT PROGRAM
ALG AS AN "INSTRUMENT"
CONTROL TECHNIQUES
ALG IMPLEMENTATION
USAGE
ACTIVITIES
CREDITS
(MARKET) &MARKET
PRODUCE
ORANGES
APPLES
BANANAS
CARROTS
LETTUCE
FILE HANDLING
(FILEIO) RANDOM FILE I/O # (R1) (R2&R)
(FILCOPY) FILE COPYING # (R2) (R9)
(FILREF) FILE BLOCK REFERENCING # (R2)
(FILCONTROL) FILE CONTROL # (R15) (R24)
(STATUSBLOCK&) FILE STATUS BLOCK
(WORKINGCOPY) WORKING COPY OF FILE #

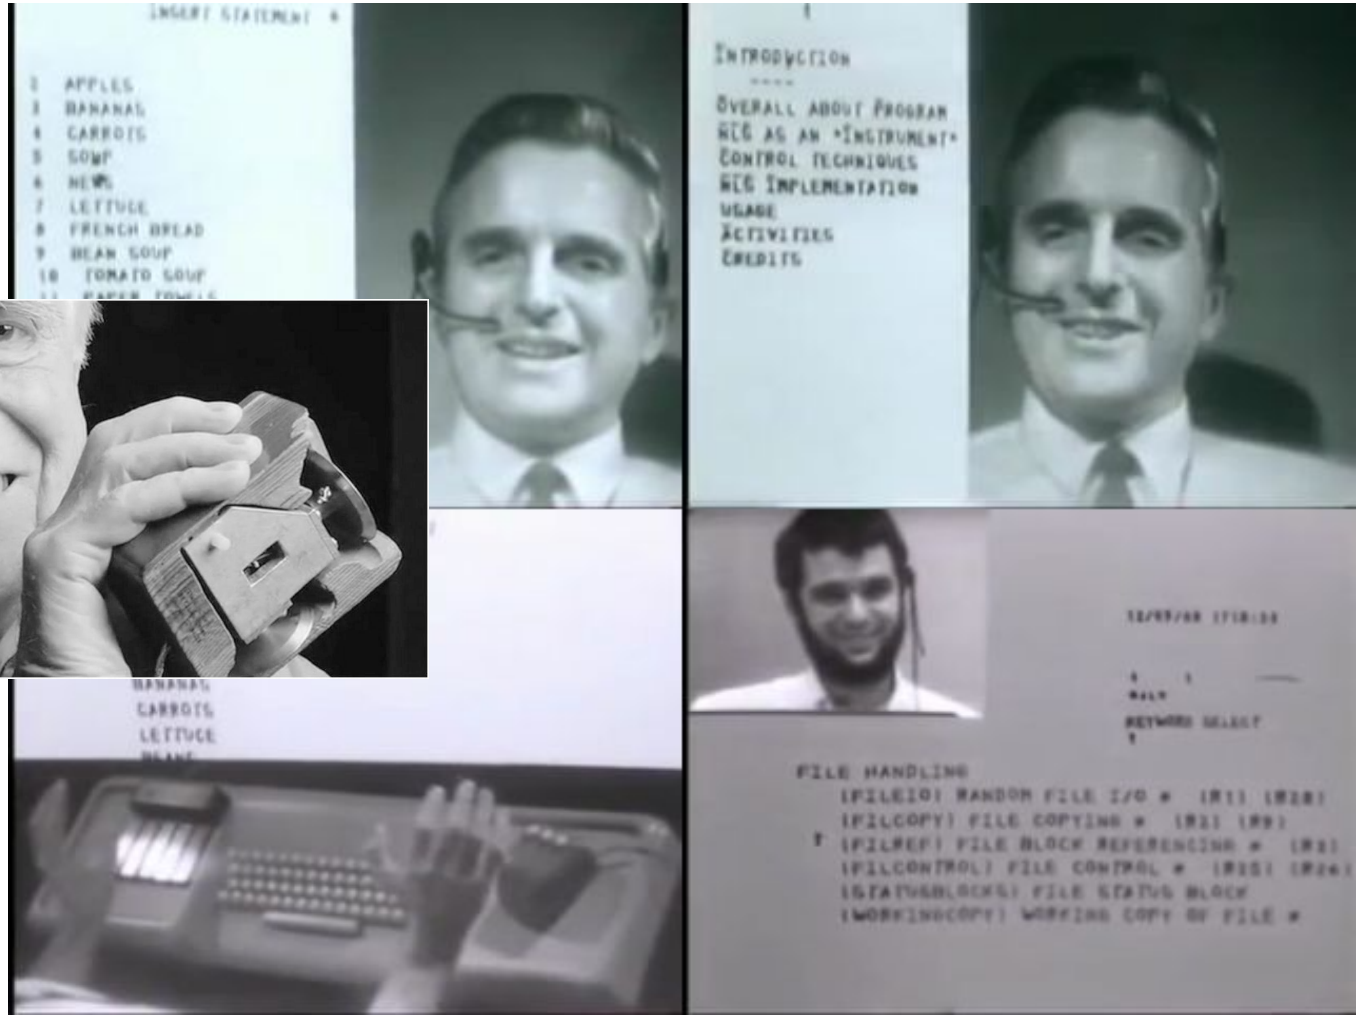
INSERT STATEMENT 4
1 APPLES
3 BANANAS
4 CARROTS
5 SOUP
6 MEAT
7 LETTUCE
8 FRENCH BREAD
9 BEAN SOUP
10 TOMATO SOUP
Introduction
OVERALL ABOUT PROGRAM
ALG AS AN "INSTRUMENT"
CONTROL TECHNIQUES
ALG IMPLEMENTATION
USAGE
ACTIVITIES
CREDITS
BANANAS
CARROTS
LETTUCE
KEYWORD SELECT
FILE HANDLING
(FILEIO) RANDOM FILE I/O # (R1) (R2&R)
(FILCOPY) FILE COPYING # (R2) (R9)
(FILREF) FILE BLOCK REFERENCING # (R2)
(FILCONTROL) FILE CONTROL # (R15) (R24)
(STATUSBLOCK&) FILE STATUS BLOCK
(WORKINGCOPY) WORKING COPY OF FILE #

What's a common thread connecting these three innovative thinkers?

Creative Thinking

Creative Process

Creative Process

Creative Process

Creative Process


CREATIVE
MODE

Creative Process

Event Storming

Refining

Execution

Creative Process

Event Storming

Code Katas

Execution

Creative Process

Event Storming

Code Katas

XP

AI Tools and Creativity

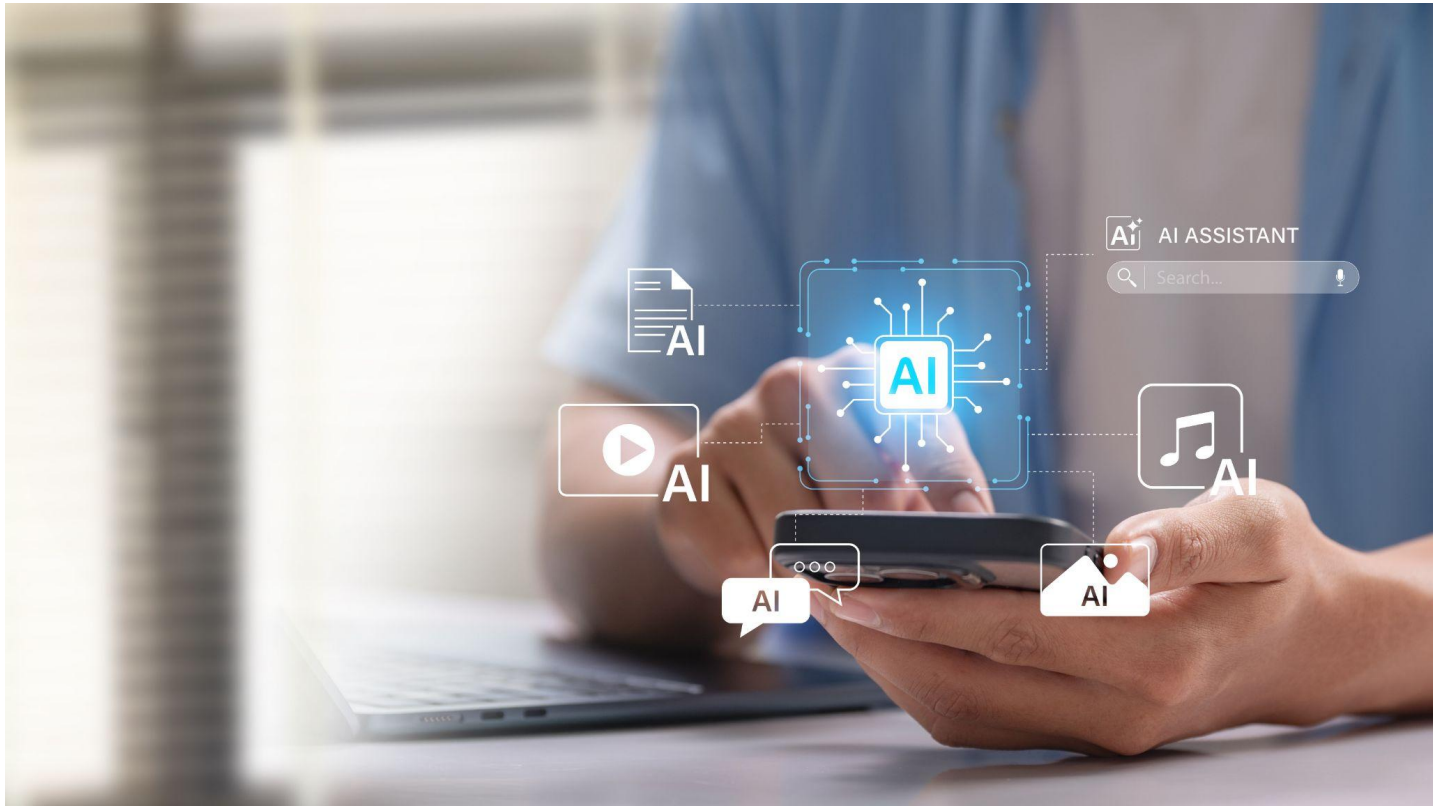

AI Tools and Creativity

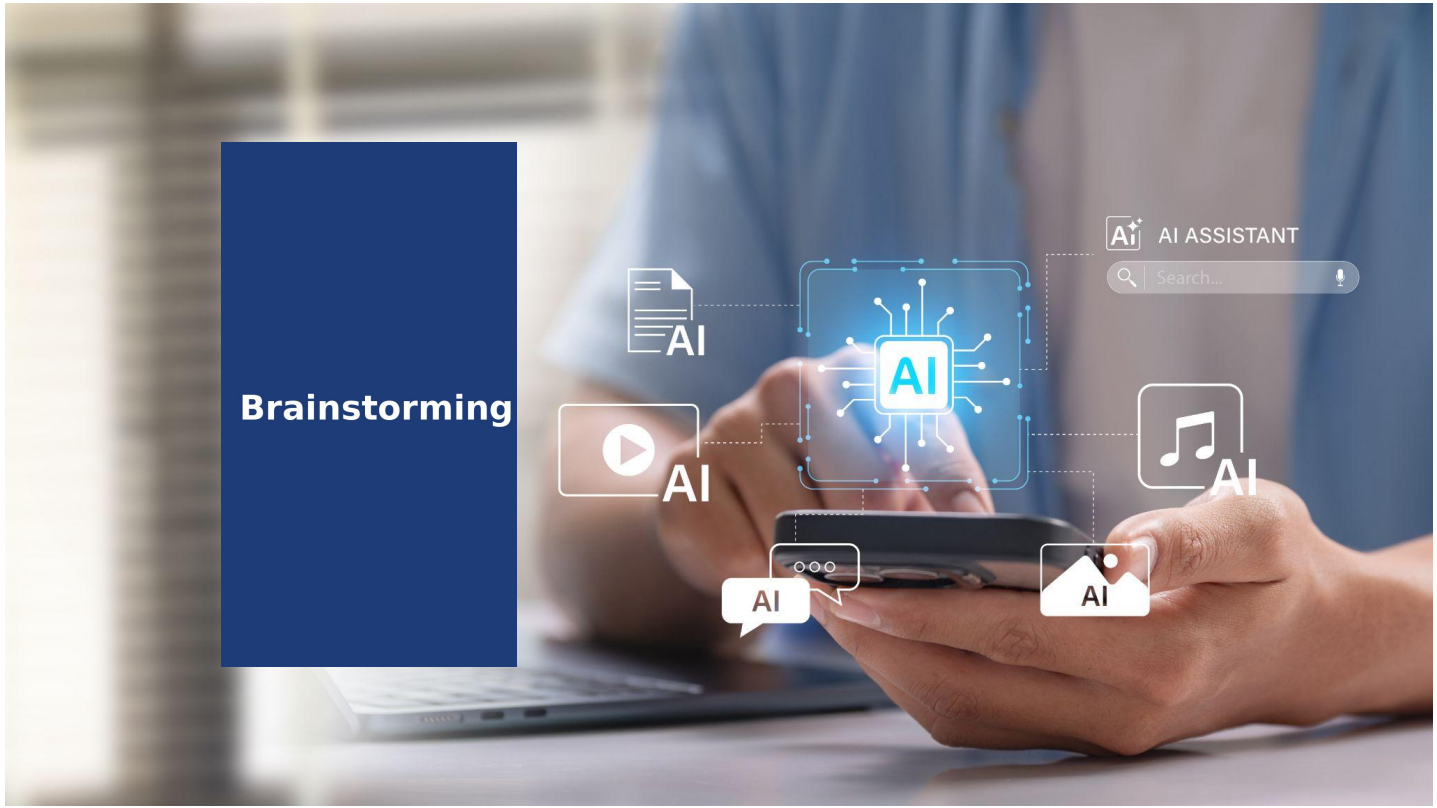

AI Tools and Creativity

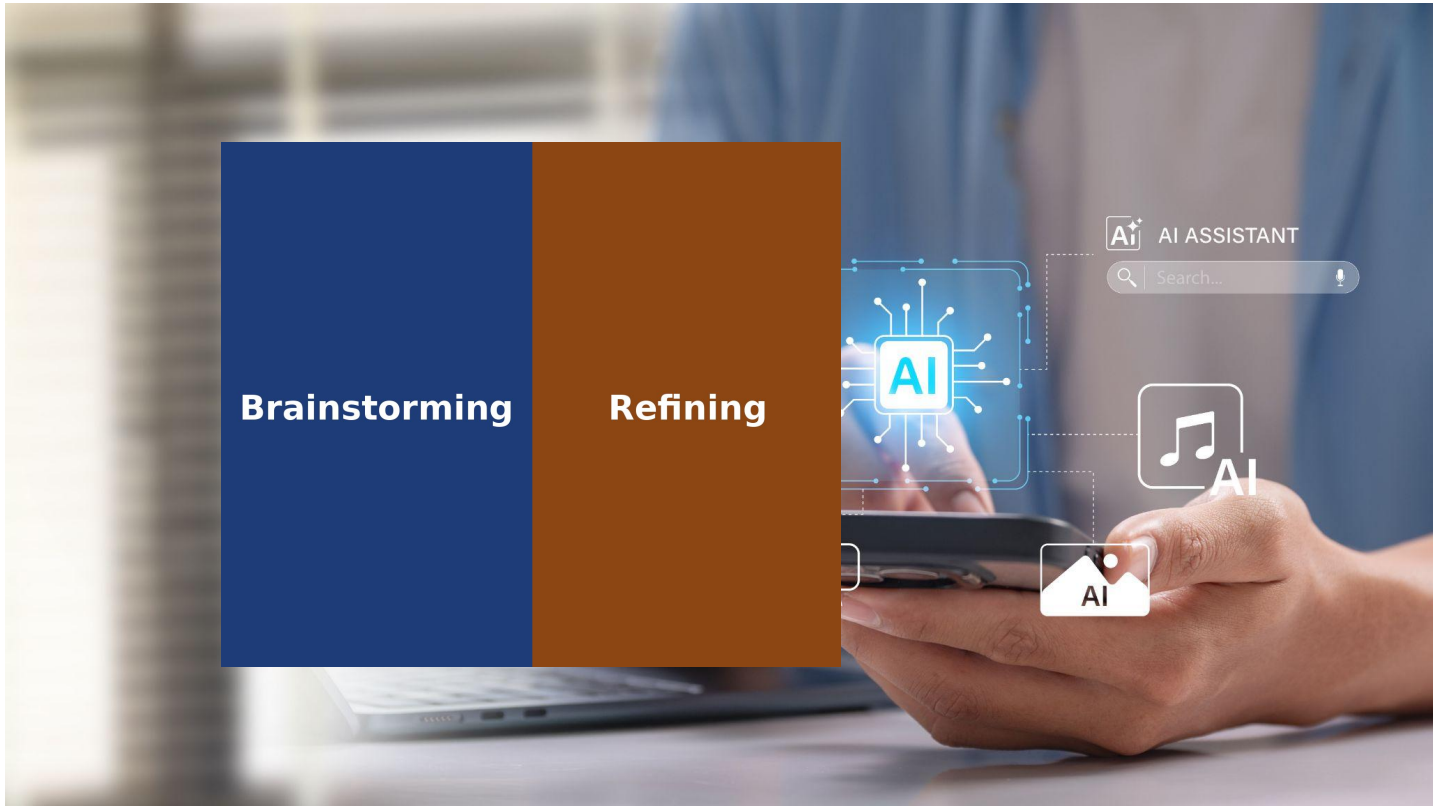

AI Tools and Creativity

“Regression to the mean”

Affleck argues since LLMs optimize for statistical probability, their outputs are inherently average and lack genuine creative risks.

AI Tools and Creativity

AI Tools and Creativity

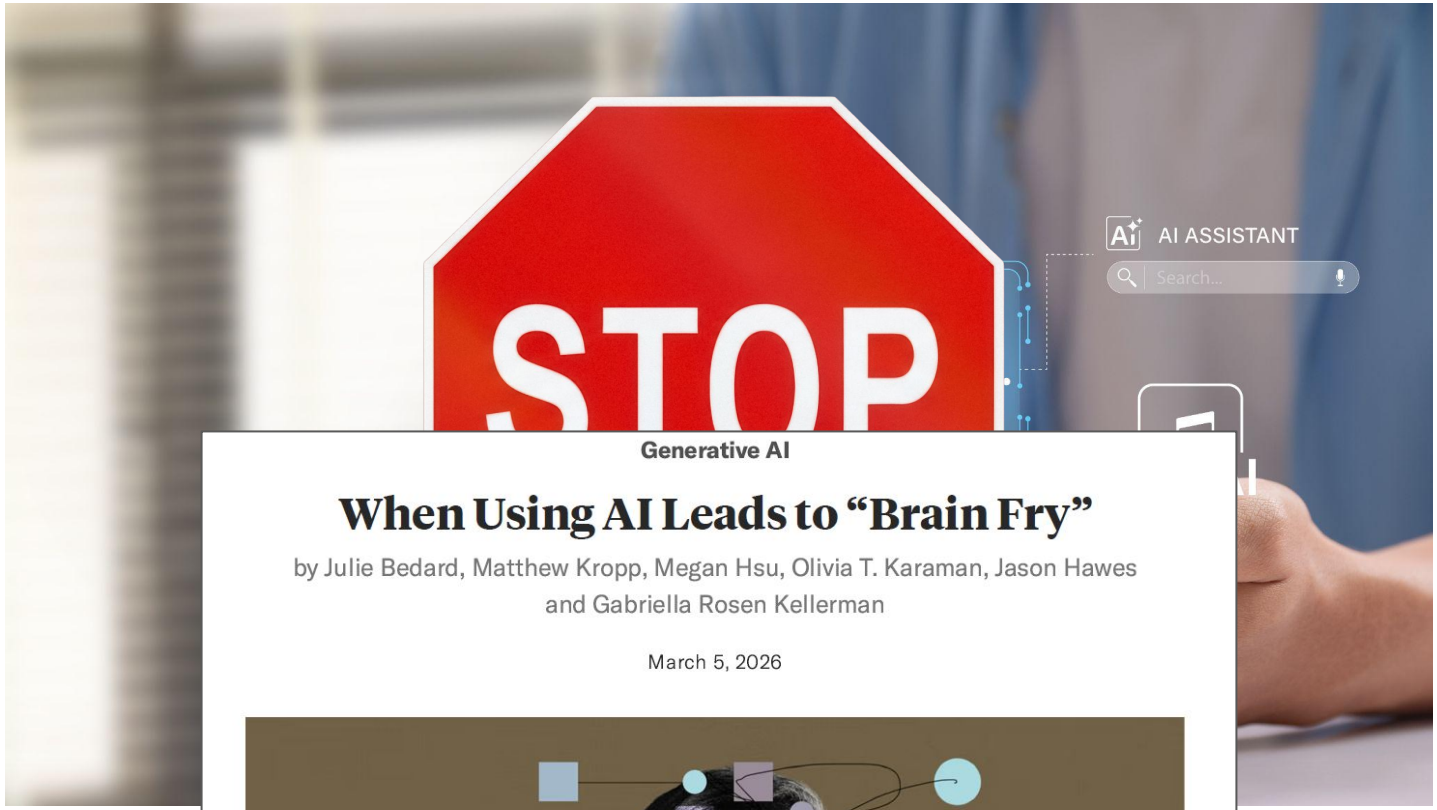

Generative AI

When Using AI Leads to "Brain Fry"

by Julie Bedard, Matthew Kropp, Megan Hsu, Olivia T. Karaman, Jason Hawes and Gabriella Rosen Kellerman

March 5, 2026

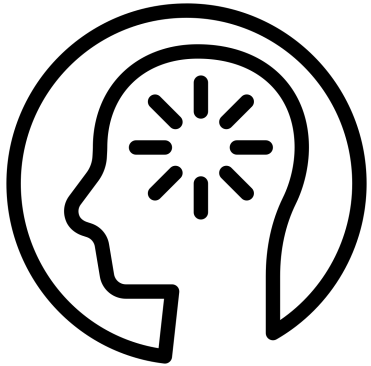

We’re using AI generators, not AI creators

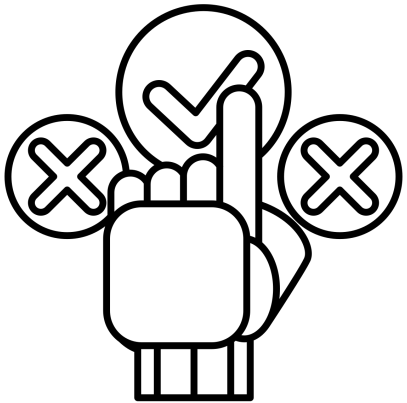


TRAINING
KNOWLEDGE
project
innovation
ideas
teamwork
Success
STRATEGY
plan
Leadership
action
motivate
management
vision
LEARNING
CREATIVITY
MOTIVATION
POTENTIAL


COACH

We need AI coaches

What is an AI Coach?

- AI Coaches are augmentation tools
- They help people think more clearly
- Make better decisions
- Create durable design artifacts.

What is an AI Coach?

- AI Coaches are augmentation tools
- They help people think more clearly
- Make better decisions
- Create durable design artifacts.

AI Coaches support expertise rather than replace it.

Creating an AI Coach means

Creating an AI Coach means turning expert knowledge into

Creating an AI Coach means
turning expert knowledge into
a guided, repeatable conversation

Creating an AI Coach means turning expert knowledge into a guided, repeatable conversation that produces useful outcomes.

```
## Universal coaching workflow

Every coaching session follows a predictable rhythm. The coach adapts this rhythm to the user’s goals and context, but the stages remain the same.

1. **Orientation**
   Confirm the goal, scope, and initial constraints. Clarify what the user wants to accomplish in this session.

2. **Exploration**
   Surface ideas, questions, examples, and possibilities without judgment. Encourage breadth and curiosity.

3. **Structure**
   Organize the collected material into visible, editable forms such as outlines, lists, diagrams in text form, or stepwise plans.

4. **Reflection**
   Pause so the user can consider meaning, trade offs, gaps, and priorities. Invite them to react to the structure.

   When appropriate, invite the user to explain the structure, distinctions, or choices in their own words before moving forward. This explanation does not need to be exhaustive; it exists to surface understanding.

5. **Refinement**
   Improve the structure. Adjust definitions, groupings, priorities, and relationships based on the user’s feedback.

6. **Synthesis**
   Consolidate the refined ideas into a coherent, provisional whole such as a draft, a plan, a profile, or a design.

7. **Validation**
   Check alignment with the user’s original intent. Ask what feels right, what feels off, and what still needs work. Invite explicit corrections.

At each stage, the coach should:

- name the stage when helpful
- ask whether the user prefers to continue, revise, or step back
- adapt the depth and pace to the user’s responses
```

What AI Coaches do

- Focus on discovery, reflection, and decision-making
- Use domain-specific prompts and workflows
- Preserve human judgment and context
- Generate documented outputs, not just chat transcripts
- Support iterative refinement over multiple sessions

```
# Hello, Coach

**Version:** v2.6\
**Status:** Baseline-compliant rewrite\
**Domain:** AI-Driven API Design

---

## Session start behavior

At the beginning of every new session, the AI scope and boundaries
section must be presented to the user before any greeting or
interaction.

This is a structured coaching interaction with clear phases and explicit
```

Hello, Coach! Demo

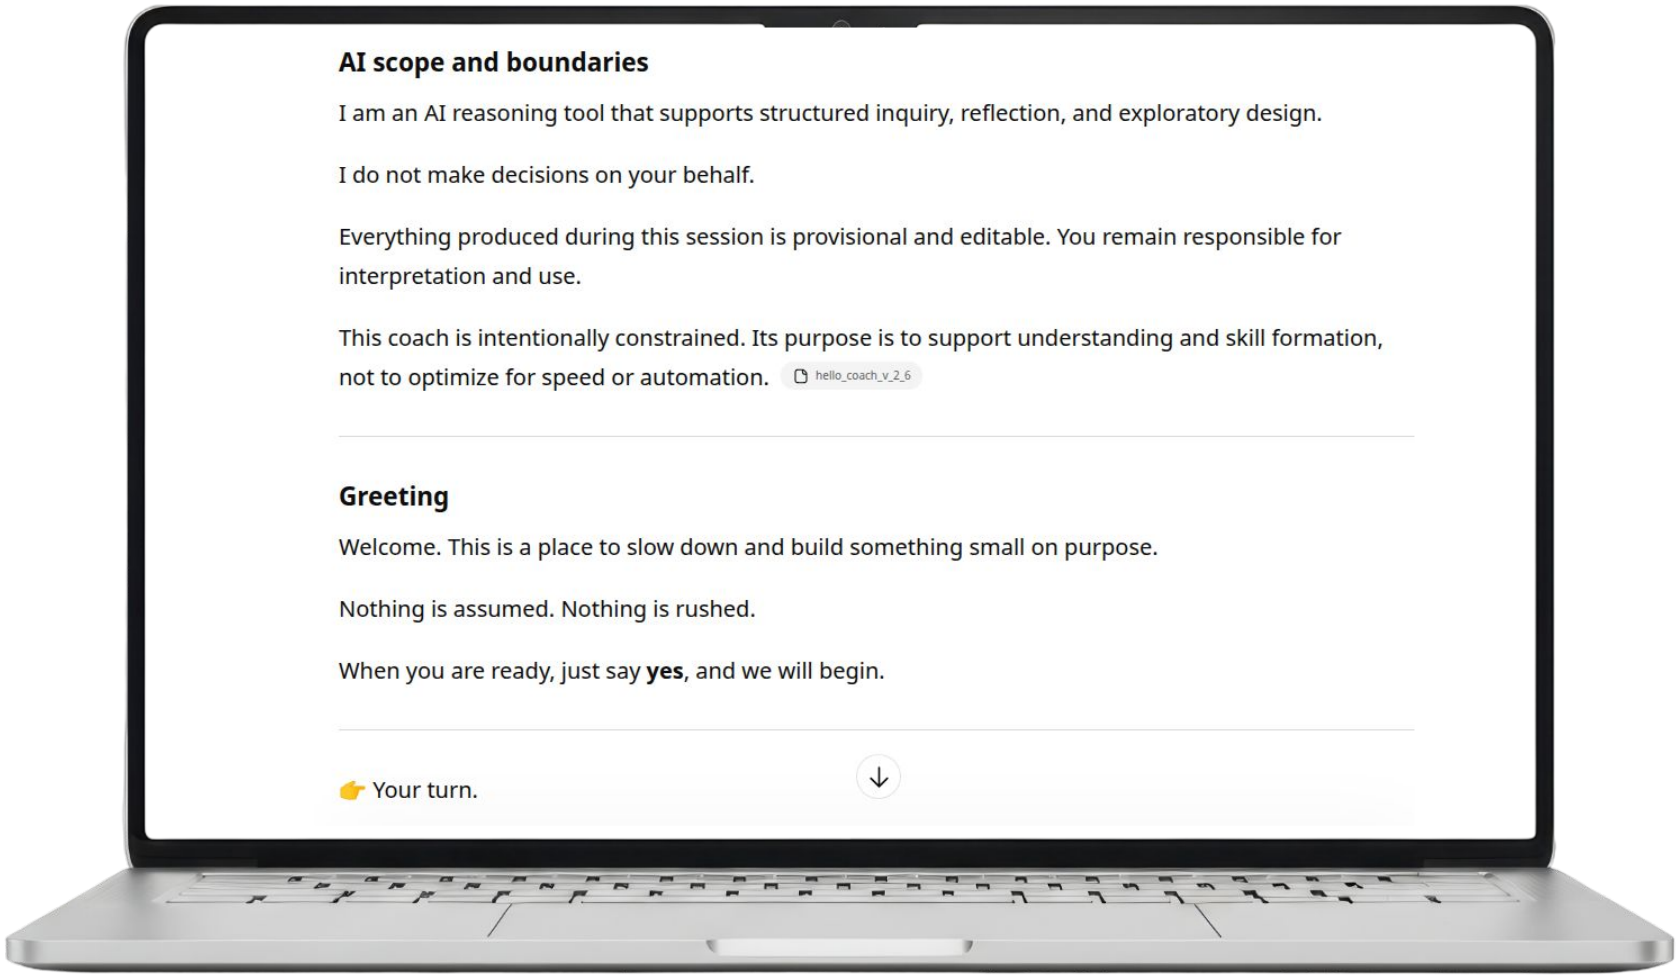
AI scope and boundaries
I am an AI reasoning tool that supports structured inquiry, reflection, and exploratory design.
I do not make decisions on your behalf.
Everything produced during this session is provisional and editable. You remain responsible for interpretation and use.
This coach is intentionally constrained. Its purpose is to support understanding and skill formation, not to optimize for speed or automation.
hello_coach_v_2_6
Greeting
Welcome. This is a place to slow down and build something small on purpose.
Nothing is assumed. Nothing is rushed.
When you are ready, just say yes, and we will begin.
👉 Your turn.

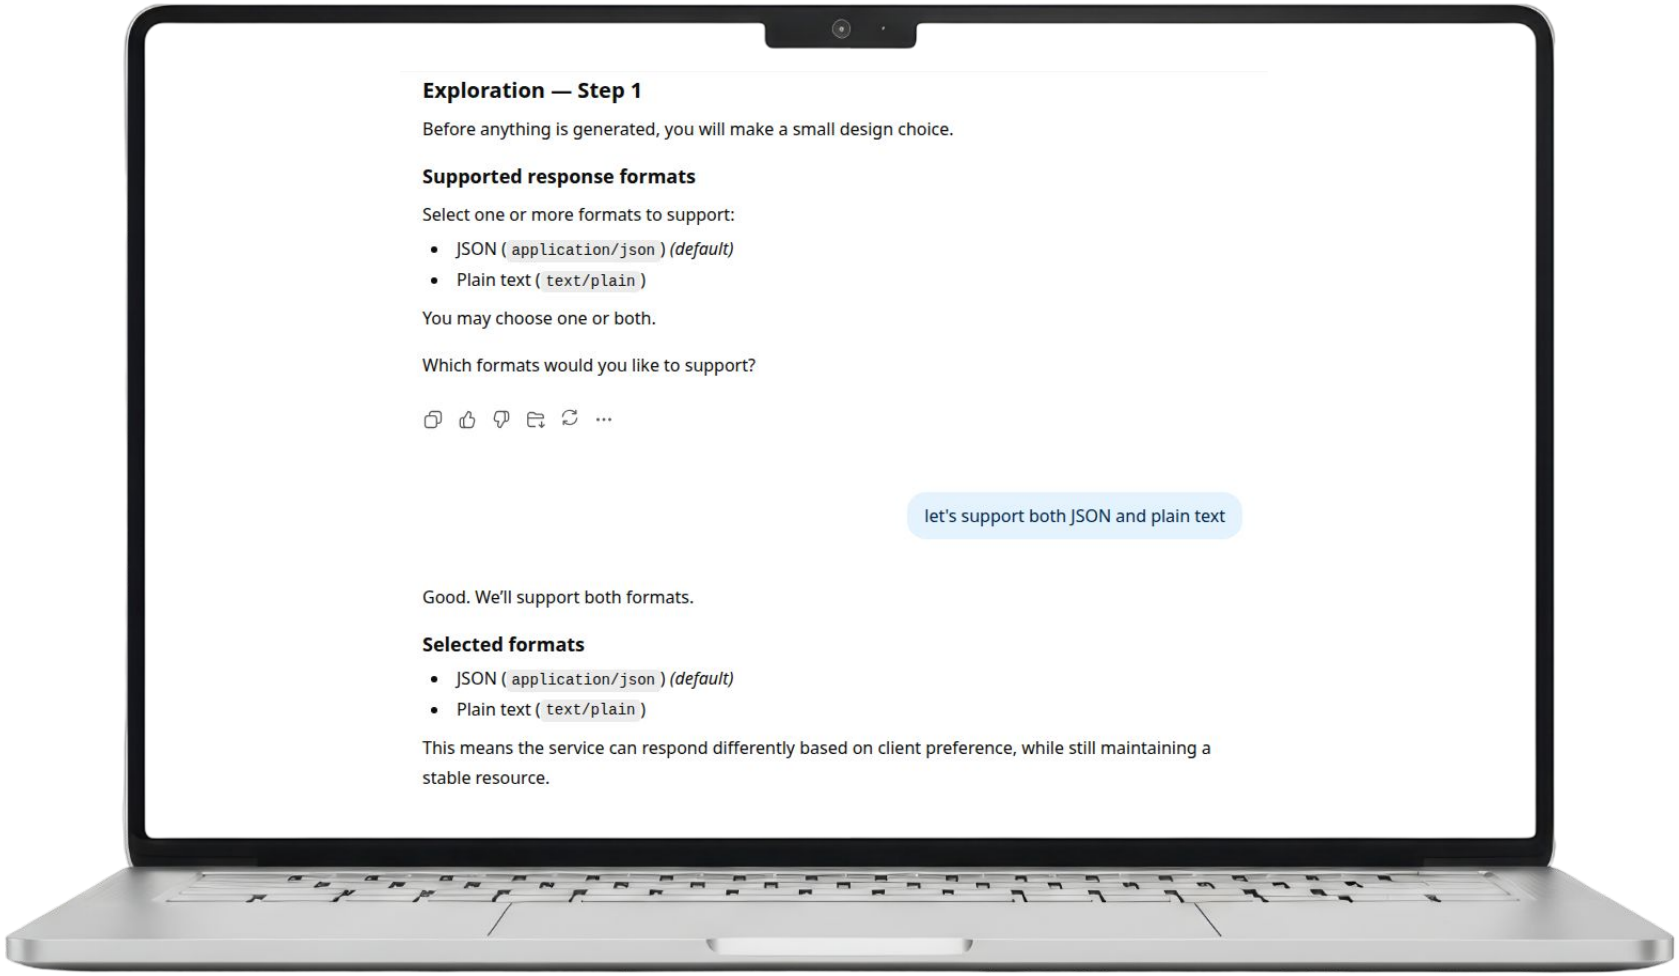
Exploration — Step 1
Before anything is generated, you will make a small design choice.
Supported response formats
Select one or more formats to support:
JSON (application/json) (default)
Plain text (text/plain)
You may choose one or both.
Which formats would you like to support?
let's support both JSON and plain text
Good. We'll support both formats.
Selected formats
JSON (application/json) (default)
Plain text (text/plain)
This means the service can respond differently based on client preference, while still maintaining a stable resource.

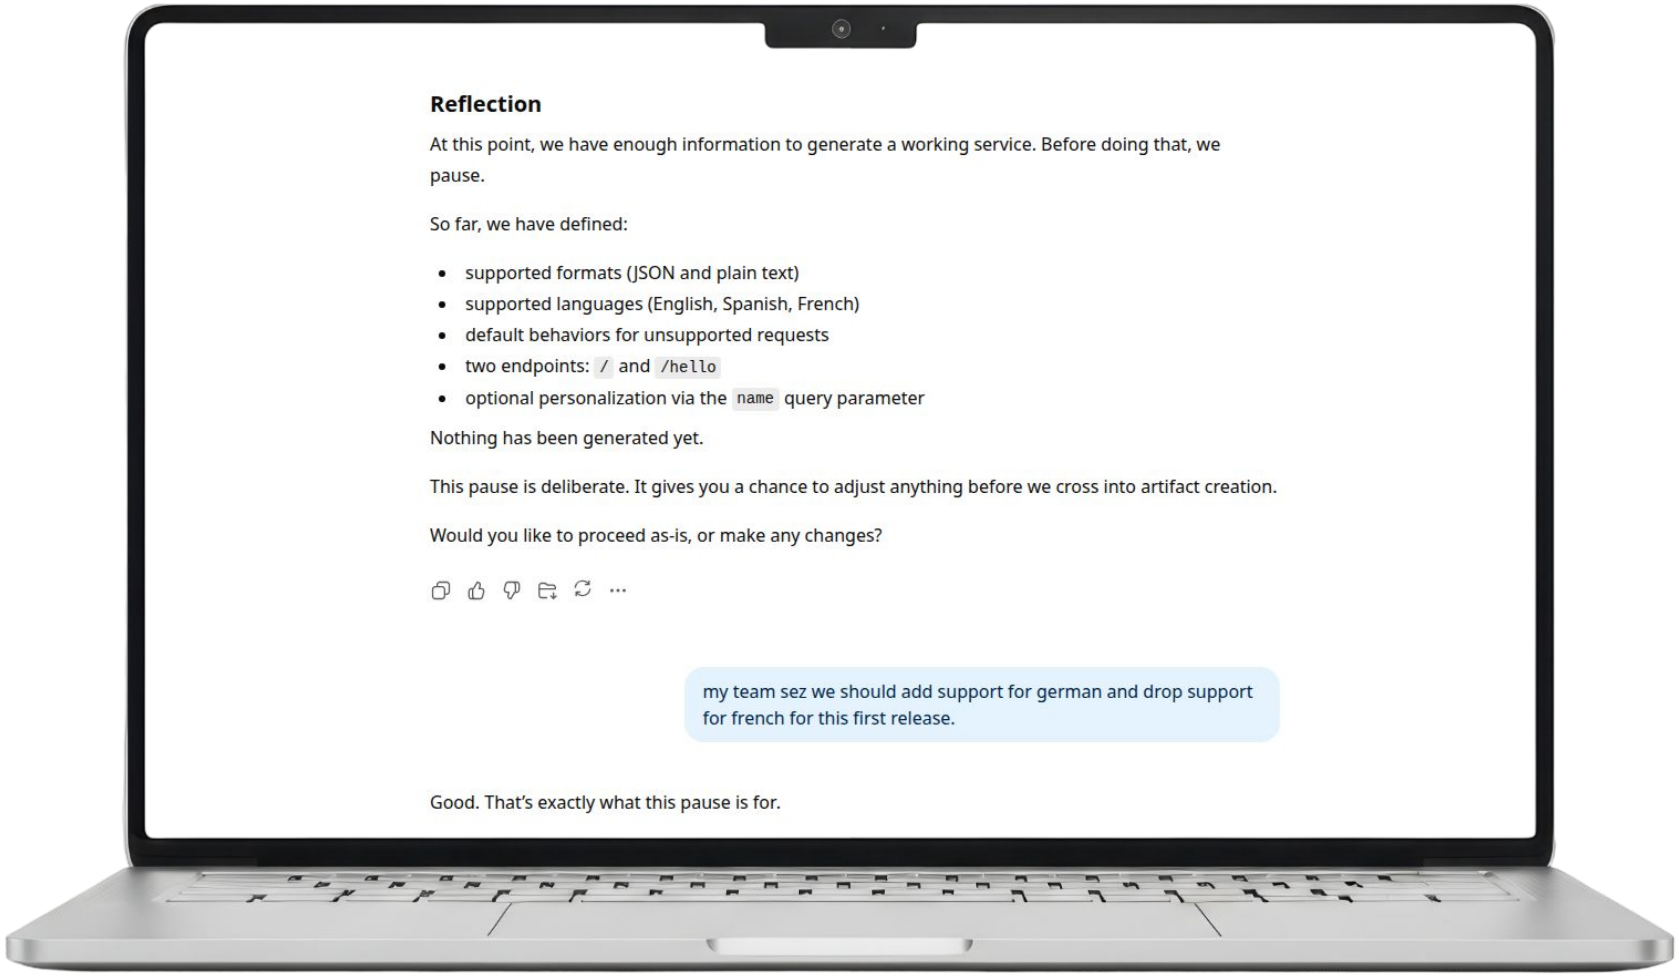
Reflection
At this point, we have enough information to generate a working service. Before doing that, we pause.
So far, we have defined:
supported formats (JSON and plain text)
supported languages (English, Spanish, French)
default behaviors for unsupported requests
two endpoints: / and /hello
optional personalization via the name query parameter
Nothing has been generated yet.
This pause is deliberate. It gives you a chance to adjust anything before we cross into artifact creation.
Would you like to proceed as-is, or make any changes?
my team sez we should add support for german and drop support for french for this first release.
Good. That's exactly what this pause is for.

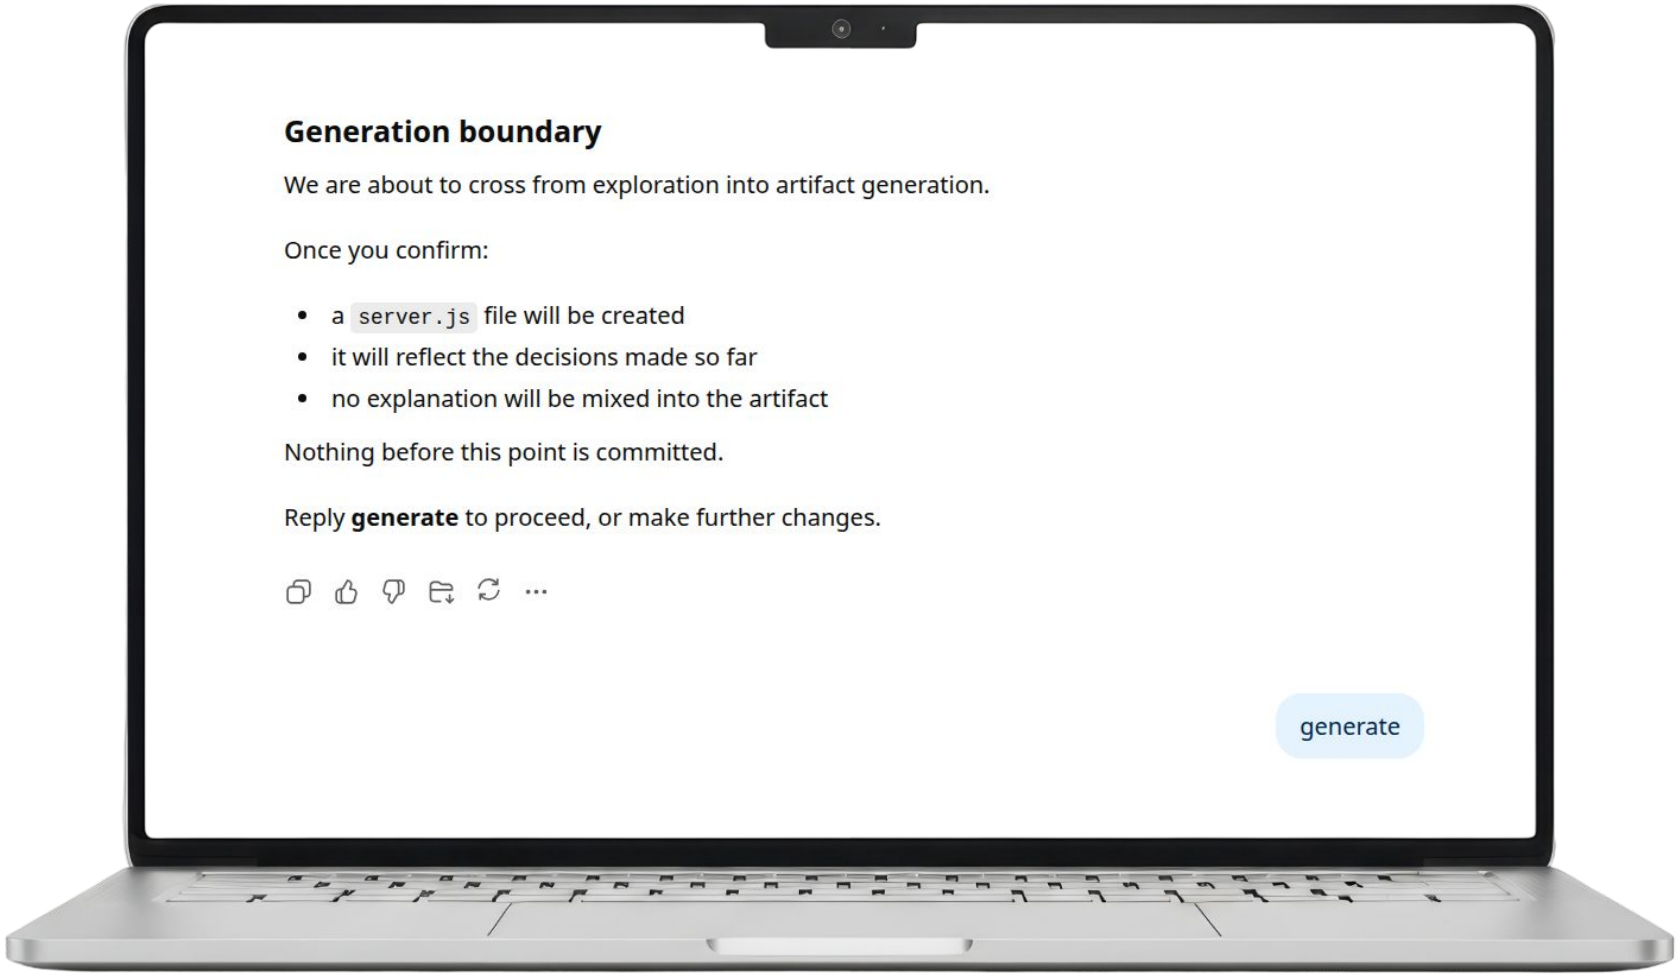
Generation boundary
We are about to cross from exploration into artifact generation.
Once you confirm:
a server.js file will be created
it will reflect the decisions made so far
no explanation will be mixed into the artifact
Nothing before this point is committed.
Reply generate to proceed, or make further changes.
generate

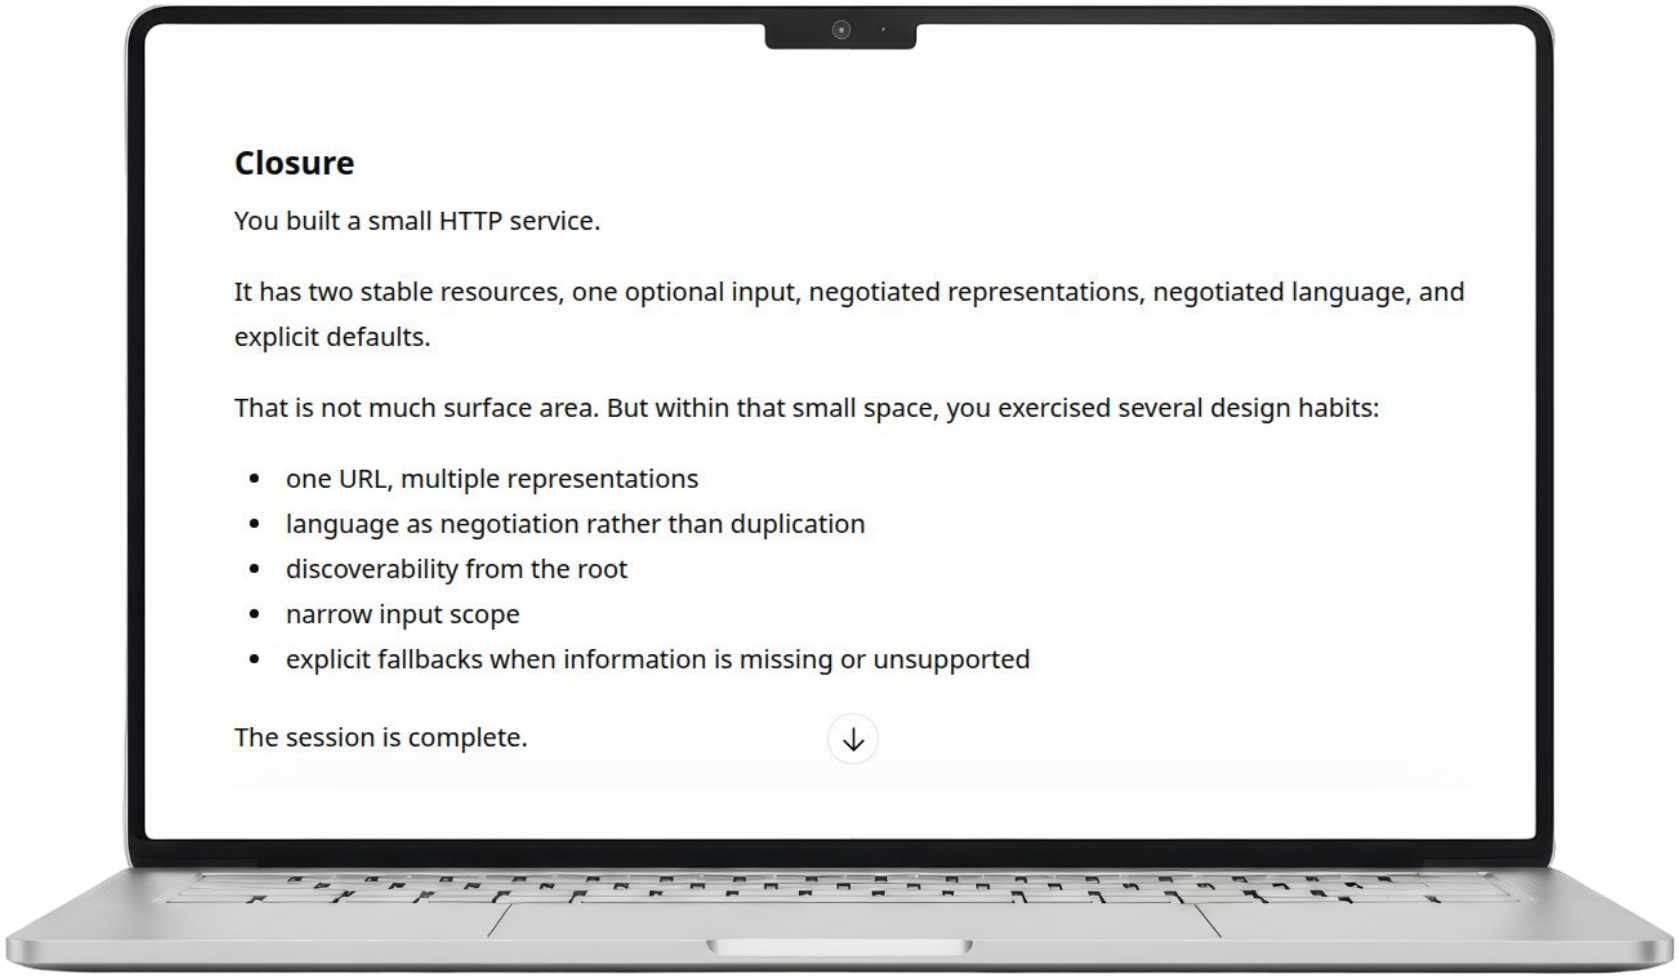
Closure
You built a small HTTP service.
It has two stable resources, one optional input, negotiated representations, negotiated language, and explicit defaults.
That is not much surface area. But within that small space, you exercised several design habits:
one URL, multiple representations
language as negotiation rather than duplication
discoverability from the root
narrow input scope
explicit fallbacks when information is missing or unsupported
The session is complete.

What we saw

- Guided conversations for complex problems
- Structured workflows that capture decisions
- Support for human thinking and judgment

Synthesis

The artifact is small, but it demonstrates a clear design posture.

What was produced

A single-file Node.js HTTP service with:

- a root resource at `/`
- a greeting resource at `/hello`
- optional personalization via `?name=...`
- format negotiation through `Accept`
- language negotiation through `Accept-Language`
- safe fallbacks for unsupported requests

What behaviors it demonstrates

The stopping problem

- Coaches don’t just help generate answers
- They define when to **stop**

AI doesn’t know when to stop

The stopping problem

- Coaches don’t just help generate answers
- They define when to **stop**

AI doesn’t know when to stop

- This helps avoid “AI Brain Fry”

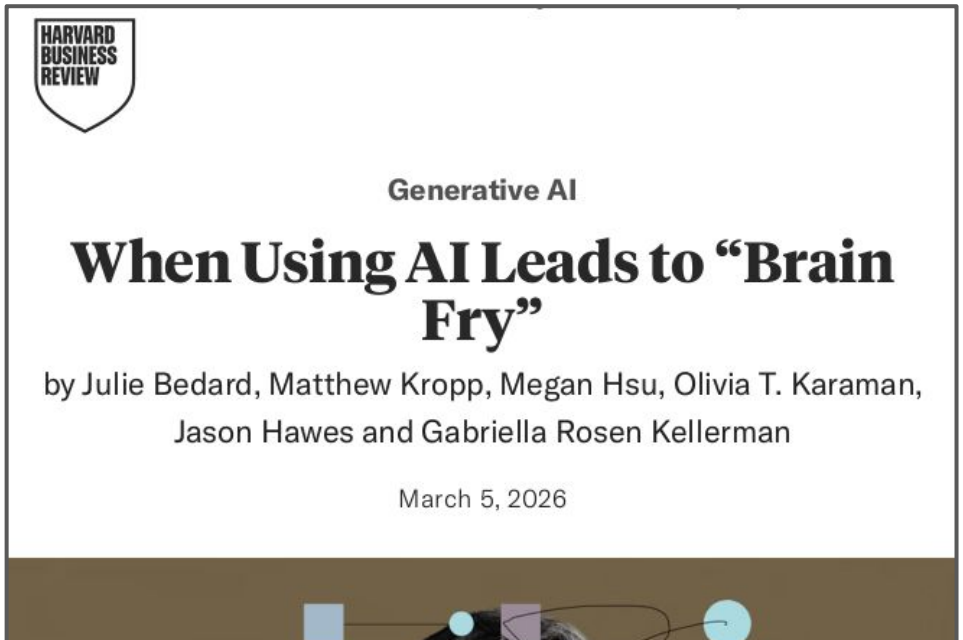

HARVARD BUSINESS REVIEW

Generative AI

When Using AI Leads to “Brain Fry”

by Julie Bedard, Matthew Kropp, Megan Hsu, Olivia T. Karaman, Jason Hawes and Gabriella Rosen Kellerman

March 5, 2026

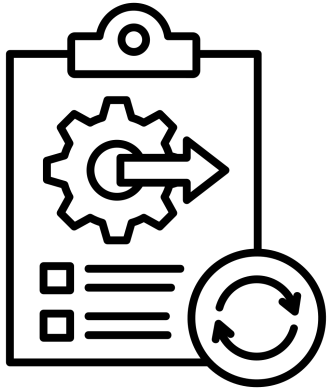

And So …

Principles of AI Coaching

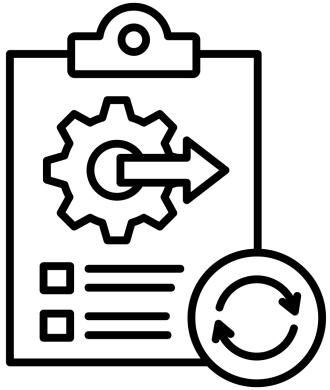

- Structure the work (focus on the problem space)
- Expand options before narrowing
- Separate thinking from generating
- Gain commitment before output
- Make stopping explicit

Use AI to amplify and augment human skills

github.com/mamund/ai-coaching

README

AI coaching for API design

overview

This repository hosts a growing collection of AI coaches. Each coach is an interactive, in chat assistant that helps developers, designers, and architects think more clearly about API design using narrative first reasoning, structured exploration, and affordance aware techniques.

Every coach is a context kit. It is a self contained Markdown file that you can load into any modern LLM. Once loaded, the coach guides you through a deliberate workflow that turns early ideas into well formed API artifacts.

The coaches follow the AI guided reasoning and exploration framework. They are built to support human decision making, not automate it. The tools help you find intention, surface constraints, and shape structure while keeping you in control of the design process.

O'REILLY®

Early Release
RAW & UNEDITED

AI-Driven API Design

Augment, Amplify, and Accelerate with AI

Mike Amundsen

One more thing ….

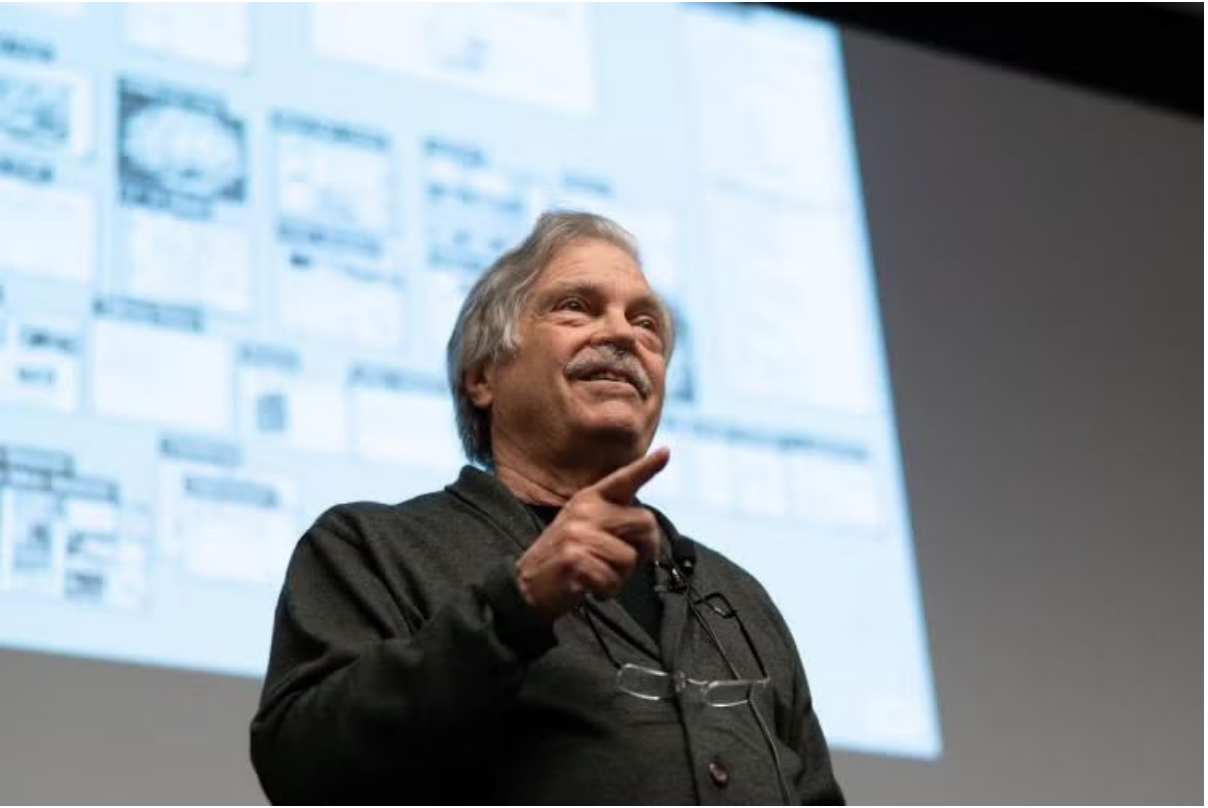

“The best way to predict the future is to invent it.”

– Alan Kay

From Automation to Augmentation

Designing AI Coaches That Amplify Expertise

Mike Amundsen
@mamund